

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
14 November 2002 (14.11.2002)


PCT

(10) International Publication Number
WO 02/091756 A1

- (51) International Patent Classification⁷: **H04Q 3/00**, (74) Agent: **HOFMAN-BANG ZACCO A/S**; Hans Bekkevolds Alle 7, DK-2900 Hellerup (DK).
H04L 29/06, 12/24
- (21) International Application Number: **PCT/EP02/04867** (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.
- (22) International Filing Date: **2 May 2002 (02.05.2002)**
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:
PA 2001 00707 7 May 2001 (07.05.2001) DK
01610120.6 22 November 2001 (22.11.2001) EP
60/334,552 3 December 2001 (03.12.2001) US
- (71) Applicant (*for all designated States except US*): **TELEFONAKTIEBOLAGET, L. M. Ericsson (publ)** [SE/SE]; s-126 25 Stockholm (SE).
- (72) Inventors: **STEENFELDT, Rico, Werni**; Ærtevej 9, 1.th., DK-2700 Brønshøj (DK). **SMITH, Henry, Gerard**; 28 Edge Hill Court, Edge Hill, SW 19 4LL London (GB).
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:
— with international search report

[Continued on next page]

(54) Title: **SERVICE TRIGGERING FRAMEWORK**



(57) Abstract: Disclosed is a method and a system for managing a plurality of services triggered by a message of a session protocol such as SIP controlling a communications session, the method comprising the steps of obtaining a number of execution rules each of which specifying a condition for invoking a service; and processing the execution rules in a predetermined order, a first execution rule causing a first service to be invoked, if the message fulfils a first condition, resulting in a first modified message; and a second execution rule causing a second service to be invoked with the first modified message as an input, if the first modified message fulfils a second condition specified by a second execution rule.

WO 02/091756 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Service triggering framework

5 This invention relates to the management of a plurality of services related to a communications session, the communications session being controlled by a session protocol providing session information about said communications session.

10 There is an increasing demand for interactive communications sessions over the Internet, such as IP telephony, multimedia sessions, video streaming, etc. Interactive communications sessions may be controlled by session protocols, such as the Session Initiation Protocol (SIP) which handles initiation, termination, and
15 modification of sessions between users. SIP is not concerned with the type of session to be initiated, i.e. with the actual content of SIP messages, but rather with the managing of a session. This includes tasks such as determining where a user to be contacted is actually
20 residing, delivering a description of the session that the user is being invited to, negotiating a common format for describing sessions, etc.

SIP is based on the request-response paradigm. For
25 example, when initiating a session, a caller sends a request addressed to the user the caller wants to call, i.e. the callee. The request message is sent to the callee, typically via a number of proxy servers responsible for routing and delivering messages. Thee
30 callee then sends a response, accepting or rejecting the invitation. The response is forwarded back through the sequence of proxy servers in the reverse order.

On top of the standard session management provided by a
35 session protocol such as SIP, additional services may be

implemented at the caller site, the callee site, or at any of the intermediate proxy servers.

Here, the term service comprises a unit of functionality which may incrementally be added to a base system and which results in an output which is perceivable by a user, such as a subscriber, an administrator, or the like. Hence, a service, e.g. call forwarding, voice mail, video conferencing, etc., is a modular extension to a base system, such as a system for managing sessions, e.g. a SIP system. The process of adding and enabling features in a base system will be called feature deployment.

During a session, features may be triggered by certain events. Triggering, i.e. the act of invoking a given application on a given event, is usually based on contractual relationships between subscribers and service providers. If more than one feature is deployed in a service network, and one or more service can be activated simultaneously for one or more users, then feature interactions occur. Here, the term feature interaction comprises the influence or modification of one feature by another. Feature interaction is an inevitable by-product of modular features. There may be desirable and undesirable feature interactions. However, it is a problem that the overall behaviour of a service network may become uncontrollable if feature interactions are not managed properly. Consequently, feature interaction is an increasing problem, as the number and complexity of service application increases with the emergence of new technologies, such as UMTS, which put heavy demands on the capabilities of service networks. These services may have been developed independently and the service providers need to be able to specify how conflicting instructions from these services are to be solved and

mediated to the default behaviour of the communications protocol used.

5 In order to avoid feature interaction, the behaviour of multiple features may be tested ad-hoc or systematically, when adding a new service to a service network, for example by testing pairs of features. If instances of feature interactions are detected during tests or after deployment, the detected problem may be fixed, e.g. by
10 re-designing one or more of the involved service applications.

The above approach requires a considerable amount of resources, in particular as the number and complexity of
15 the services increases. Hence, it is a problem that the above prior art does not scale well with the number and complexity of service applications.

According to the invention, the above and other problems
20 are solved by a method of managing a plurality of services triggered by a message of a session protocol controlling a communications session,

the method comprising the steps of
25 obtaining a number of execution rules each of which specifying a condition for invoking a service;

processing the execution rules in a predetermined order,
30 a first execution rule causing a first service to be invoked, if the message fulfils a first condition, resulting in a first modified message; and a second execution rule causing a second service to be invoked with the first modified message as an input, if the first
35 modified message fulfils a second condition.

Consequently, the invocation of services triggered during a communications session is controlled by a number of execution rules which are processed in a predetermined order, thereby controlling the order of services to be invoked. Therefore, a mechanism for triggering applications based on the service execution rules is provided which avoids an arbitrary overall behaviour due to an uncontrolled order of execution. Furthermore, by editing the execution rules, different deployment strategies may easily be implemented, thereby providing a flexible, fine-grained service deployment infrastructure which provides great flexibility in ordering the chain of services allowing to optimise the performance of the service network.

Hence, according to the invention, a flexible service deployment infrastructure is provided which allows to systematically avoid feature interactions when managing a large number of complex services, e.g. when adding, removing, suspending, re-activating, or re-locating services within a service network.

A standardised framework for defining service execution is provided which allows the distribution of the service management problem between independent stakeholders, thereby providing a method which is scalable with the number of services.

An execution rule includes one or more conditions for performing one or more actions, such as invoking a service application.

The term communications session comprises communications sessions between users of a communications network such as a TCP/IP network, a local area network, a wide area network, the Internet, or the like. Examples of

communications sessions include Voice-over-Internet, IP telephony, video conferencing, video streaming, etc.

5 The term session protocol comprises a protocol controlling the communication session, and in particular the initiation, termination and modification of sessions. Preferably, the session protocol is based on request/response messages transmitted via the nodes of the communications network between the participating
10 users of a communications session. An example of such a session protocol is the Session Initiation Protocol (SIP). Other examples include the H.323 protocol suite, MGCP and related protocols such as IPDC, SGCP, H.248 etc.

15 The term service comprises a unit of functionality which may incrementally be added to a base system. A service may offer services to subscribers, but may also offer other services, like administrative tasks to the network administrator. Examples of such services include call
20 forwarding, call waiting, voice mail, statistics functions, call back, video conferencing, video on demand, anonymizer services, auto reply, etc. A service may be implemented using a variety of technologies, e.g. OSA, Java, CGI, Perl, C++, CPL, XML, etc.

25 In the context of the SIP protocol, a service is an application or a number of applications executed locally on a SIP server, e.g. as a CGI-Script or a CPL-Script, or remotely on an application server contacted by the SIP
30 server. In the latter case, the service may be accessed and invoked using some standard naming convention, i.e. using SIP Request-URIs. Alternatively, the services may actually register themselves at the SIP server using e.g. a 3GPP OSA API framework. SIP services can be grouped
35 into originating and terminating services, i.e. those that are associated the caller and the callee.

Services may be triggered on conditions in a message header, a message body, the SDP, or the like.

5 The functionality of the SIP Server that is offered to service applications is termed service features, such as access to some API, e.g. a server side OSA API, to statistical functions, or the like. Service features may further be service applications which register at the SIP
10 Server and subsequently offer their service to be used by other service applications.

It is a further advantage of the invention that it is independent of the technology used for the implementation
15 of the services, the signalling protocol, and the platform. Hence, a network operator does not need to know in advance which types of services that will be deployed in the service network, thereby providing a robust and extendible service deployment infrastructure.

20 As the number of stakeholders that may want to register their own services could be very large, there is a need for scalability. Furthermore, the number of subscribers associated with a domain may be very large, rendering the
25 issue of scalability critical. Services may be triggered by many different types of events and invoked based on a plurality of different conditions, such as matching source and destination addresses, time-dependant, or some other pre-condition. Furthermore, non-SIP related
30 services may be invoked on SIP events, e.g. if certain conditions are fulfilled at a given point in time. Different service technologies, such as CPL, may be used. SIP related services may be invoked on non-SIP related events, e.g. HTTP events. There may be tens of thousands
35 of services that may be offered to tens of millions of subscribers, from tens of thousands 3rd party service

providers. Consequently, the task of managing services and service interactions is a task which easily gets too big and complicated for one administrator to manage.

5 According to a preferred embodiment of the invention, said number of execution rules is grouped into a number of rule modules, each rule module including a number of execution rules; and the method further comprises the steps of

- 10 - processing a first one of said number of rule modules resulting in a first accumulatively modified message; and
- invoking processing a second one of said number of rule modules providing the first accumulatively
15 modified message as an input.

Consequently, by grouping the execution rules into rule modules, i.e. groups of execution rules, the problem of feature interaction is split into the problem of feature interaction between features invoked within the same rule
20 module and interactions between rule modules. Consequently, it is an advantage of the invention, that a method of managing feature interaction is provided which scales with the number of services. In particular, when different services are provided by different service
25 providers, e.g. different companies or different organisational entities within a company, a single provider may not have access to or may not be able to test all of the services provided. Hence, it is an advantage of the invention that the task of analysing
30 feature interaction analysis may be split up according to rule modules and distributed to different providers. This further implies that the costs of service management may be delegated to independent parties as the number of subscribers and services grows.

The stakeholders that may want to upload/register and administer services on a SIP server could be the owner, provider or administrator of the SIP server, network operators, etc. It could also be different types of
5 retailers, e.g. virtual telecom operators, Internet Service Providers, etc. It may also be different types of service providers, such as application service providers, content service providers, service/feature providers. Also private organizations, enterprises and subscribers
10 may be possible stakeholders.

Consequently, according to the invention, a flexible, extensible and scalable management of contractual relationships between the stakeholders is achieved,
15 including the management of charging, settlements, policies and security.

It is a further advantage of the invention, that it provides a modular structure of execution rules, thereby
20 enabling the embedding of service profiles for users, user groups, subscribers, etc. It is a further advantage of the modularity that it enables reuse of rule modules, thereby further facilitating the maintenance of the service environment.

25 When each rule module has associated with it a priority indicative of an order of processing of said number of rule modules, the order of rule module execution is determined by a simple parameter, thereby providing easy
30 and transparent control over the order of execution of rule modules.

When each rule module corresponds to a rule module owner authorised to edit the rule module, the administrative
35 authority for a rule module may easily be established, thereby further facilitating the delegation of

administrative authority, such as editing rule modules, feature interaction analysis within a rule module, etc.

According to a further preferred embodiment of the invention, the first rule module has assigned to it a privilege indicative of an authority to alter a lock flag related to a predetermined part of the accumulatively modified message and specifying whether said predetermined part of the accumulatively modified message may be modified by services invoked from at least the second rule module. Consequently, a mechanism is provided for explicitly allowing or preventing alterations of individual attributes or groups of attributes of the messages by subsequent services. Hence, an efficient tool is provided for avoiding feature interactions due to the changing of the context of one service by another service. This type of feature interaction will be referred to as violation of feature assumption which may cause ambiguous or conflicting behaviour of services. As the authority to lock and/or unlock certain attributes is linked to predetermined privileges assigned to rule modules, the network operator may detect misuse of privileges and thereby prevent unauthorised behaviour, thereby increasing the security of the method. Privilege violation may be detected at run-time and resolved, e.g. by notification and or de-activation of services.

According to a yet further preferred embodiment of the invention, the step of invoking processing the second rule module further comprises the step of setting said lock flag to prevent modification of the predetermined part of the accumulatively modified message by services invoked from the second rule module, unless the lock flag was marked unset by the first rule module. Consequently, by default, the message attributes are locked for subsequent changes when the control is transferred from

one rule module to another. Hence, the second rule module may only change message attributes which the first rule module has explicitly marked as being unlocked, thereby further improving the control of possible feature interactions between rule modules and confining feature analysis task to within the individual rule modules. This results in a further improved scalability.

When the step of obtaining a number of execution rules further comprises the step of detecting a predetermined contractual relationship based on header information of the message; and selecting a number of rule modules based on said detected contractual relationship, a modular and scalable method is provided which enables the support for operators to host 3rd party services and/user defined services. As contractual relationships are detected on the basis of header information, those 3rd party or user-defined services which are relevant for a given message may be identified on the basis of the detected contractual relationships.

According to another preferred embodiment of the invention, the step of processing the first rule module further comprises the step of invoking a predetermined third rule module. Consequently, a rule module may invoke other rule modules, thereby providing a powerful tool for creating a hierarchy of rule modules and further improving the scalability and flexibility of the deployment infrastructure. It is an advantage of the invention that it allows the delegation from one rule module owned by one party to another rule module owned by another party. One party may invoke applications in an order which is deemed appropriate by that party and then pass control to another party which then may invoke different applications. The first party may subsequently regain control to check the output from the second party.

When this delegation is done, the first party may decide which messages properties in the forwarded message subsequent applications can change by indicating which message properties cannot be overwritten for an action.

5

It is an advantage of the invention that it allows a hierarchical delegation of administrative authority based on the ability to trigger rule modules from within other rule modules and lock message properties.

10

When the first and second rule modules are related to respective first and second access control lists specifying access rights to the corresponding first or second rule module, the violation of access rights may be detected and resolved, thereby further increasing the security of the method.

15

When the first and second rule modules comprise respective first and second scripts in a predetermined mark-up language, a simple language for writing rule modules is provided. Hence it is easy for an administrator to understand how to express rules and what the meaning of the rules is, thereby providing a simple mechanism for extending the basic definition of the rule language with proprietary functions. An example of such a language definition is XML.

20

25

It is a further advantage of the invention that it provides an extensible framework, i.e. it is easy to add protocols, service technologies and message properties, e.g. if a new method is added to SIP.

30

It is a further advantage of the invention that it provides a scalable framework, i.e. the same way of expressing rules is possible on large ISP networks and small end-user devices, e.g. 3G cell-phones.

35

According to another preferred embodiment of the invention, the message comprises a first and a second set of attributes; the execution rules are grouped into at least a first and a second processing class of execution rules according to corresponding constraints, where the second processing class is restricted to only modify attributes of the second set of attributes; and the step of processing the execution rules further comprises the step of processing the execution rules of the first processing class before processing any execution rule of the second processing class. Consequently, the services are divided into groups with predetermined behaviour in terms of which parts of the signalling messages they update. Hence, the services may rely on that the first set of message attributes will not be altered after the execution of the first processing class. This provides a further mechanism for splitting up the task of feature interaction analysis into manageable subtasks. In the following, the processing classes will also be referred to as processing points. The sets of attributes may comprise message header information and message body comprising the actual session content. The attributes may be further split up into different types of header information, e.g. signalling attributes etc.

It is a further advantage of the invention that it provides means for specifying how applications are allowed to interact.

Preferably, the processing classes are processed in a predetermined order for requests and responses and applied to originating, terminating and "forwarded by" services.

According to a further preferred embodiment of the invention, the message comprises a first and a second set of attributes; the execution rules are grouped into at least a first and a second processing class of execution rules according to corresponding constraints, where the second processing class is restricted to only modify attributes of the second set of attributes; and the method further comprises the step of repeating the steps of processing the first rule module and invoking processing the second rule module, where in each repetition the processing of the first and second rule modules is limited to execution rules of a corresponding processing class, and where each repetition results in a corresponding accumulatively modified message which is used as an input for a subsequent repetition. Consequently, a combination of the division onto rule modules with the concept of processing classes is provided, thereby providing a fine-grained framework for dividing the services according to administrative ownership and constraints imposed on the services.

When the processing classes are defined separately for execution rules triggered by requests and responses of the session protocol, the division of services into processing classes is further split up according to the type of message, thereby providing a more fine-grained splitting.

According to a preferred embodiment, the processing classes correspond to predetermined locations in a round trip message flow according to the session protocol, thereby simplifying the analysis of feature interaction.

Preferably, the processing classes include a first processing class of execution rules which impact signalling properties of the message, a second processing

class of execution rules which impact non-signalling message body content of the message, and a third processing class of execution rules which neither impact the signalling properties nor the non-signalling message
5 body content of the message. Here, the term signalling properties comprises SIP and SDP message properties that can be matched in a rule module condition to invoke a service.

10 When a resulting modified message is generated when all execution rules of the first and second processing classes are processed, the efficiency of the method is further increased, as responses may be returned without having to wait for services of the third processing
15 class.

According to another preferred embodiment of the invention, invoking the first service further results in a second modified message; and the method further
20 comprises the steps of processing subsequent execution rules with the first modified message as an input; and processing subsequent execution rules with the second modified message as an input. Hence, as a service may return different outputs, each of which is an input to a
25 chain of subsequent services, a tree of cascaded chains of services may be implemented.

According to yet another preferred embodiment of the invention, the method further comprises the steps of
30 - storing information about which services are executed and information about which order the services are executed in;
- receiving from the first service a request for returning a notification to the first service, if a
35 predetermined event occurs;

- storing the request in relation to the stored information; and
- upon occurrence of the event, notifying the first service according to the stored information.

5 Hence, an efficient mechanism is provided for requesting notification of future events by a service application, thereby enabling monitoring applications, etc.

10 When the execution modules comprise computer-readable scripts, and the predetermined order of processing the execution rules is determined by the order of execution rules in said scripts, a simple mechanism is provided for controlling the order of service execution and the order of notification regarding future events by an
15 administrator of a rule module.

Preferably, the cascaded order of services is determined by the order of execution rules within a rule module and the order of rule modules.

20

When the execution rules are adapted to be dynamically updated, i.e. during operation of the service network, a real-time service management is provided.

25 The invention further relates to a data processing system comprising a service execution environment module adapted to invoke a plurality of services triggered by a message of a session protocol controlling a communications session;

30

characterised in that

the data processing system further comprises a storage medium adapted to store a plurality of execution rules
35 each of which specifying a condition for invoking a service; and

the service execution environment module comprises a rule engine module adapted to

- retrieve a number of execution rules; and
- 5 - process the execution rules in a predetermined order, a first execution rule causing a first service to be invoked, if the message fulfils a first condition, resulting in a first modified message; and a second execution rule causing a second service to be invoked
- 10 with the first modified message as an input, if the first modified message fulfils a second condition.

The invention further relates to, in a data processing system, a service execution environment module adapted to

15 invoke a plurality of services triggered by a message of a session protocol controlling a communications session;

characterised in that

20 the service execution environment module comprises a rule engine module adapted to

- retrieve a number of execution rules each of which specifying a condition for invoking a service; and
- process the execution rules in a predetermined order,
- 25 a first execution rule causing a first service to be invoked, if the message fulfils a first condition, resulting in a first modified message; and a second execution rule causing a second service to be invoked with the first modified message as an input, if the
- 30 first modified message fulfils a second condition.

The invention further relates to software program comprising code means adapted to perform, when executed on a data processing system, the steps of the method

35 described above and in the following.

The software program may be embodied on a computer-readable medium. The term computer-readable medium may include magnetic tape, optical disc, digital video disk (DVD), compact disc (CD or CD-ROM), mini-disc, hard disk, 5 floppy disk, ferro-electric memory, electrically erasable programmable read only memory (EEPROM), flash memory, EPROM, read only memory (ROM), static random access memory (SRAM), dynamic random access memory (DRAM), synchronous dynamic random access memory (SDRAM), 10 ferromagnetic memory, optical storage, charge coupled devices, smart cards, PCMCIA card, etc.

The invention further relates to a method of managing the deployment of a service in a service network, the method 15 comprising the steps of

- specifying in a computer-readable script a number of privileges and rights to be granted to the service during operation;
- analysing potential causes of feature interaction;
- 20 - specifying a deployment strategy for said service as a set of execution rules stored as a computer readable rule module script.

The invention further relates to a data record comprising 25 a rule module for use in the method described above and in the following.

These and other aspects of the invention will be apparent from and elucidated with reference to the embodiments 30 and with reference to the drawings, in which:

- fig. 1 illustrates different types of feature interactions in a SIP service network;
- fig. 2 illustrates the network elements involved in a 35 service environment SIP server architecture according to an embodiment of the invention;

fig. 3 schematically illustrates a system architecture of a SIP server which supports a service execution rule mechanism according to an embodiment of the invention;
fig. 4 illustrates the structure of a rule module
5 according to an embodiment of the invention;

fig. 5 illustrates the grouping of services into constrained sets of services according to an embodiment of the invention;

fig. 6 illustrates the grouping of services into
10 constrained sets of services corresponding to locations in the round trip SIP message flow according to a preferred embodiment of the invention;

fig 7 illustrates the processing flow between services belonging to different processing points according to an
15 embodiment of the invention;

fig. 8 illustrates the grouping of services into rule modules corresponding to administrative authority according to an embodiment of the invention;

Fig. 9 illustrates an embodiment of the invention where
20 services are grouped both according to processing points and according to administrative authority;

fig. 10 illustrates the processing mechanism of the embodiment of fig. 9 in the case of multiple rule modules and multiple processing points;

25 fig. 11 illustrates another example of the processing rules described in connection with fig. 9;

fig. 12 illustrates a hierarchical rule module processing according to a preferred embodiment of the invention;

fig. 13 shows an example of the flow of SIP message event
30 contexts and instruction sets according to an embodiment of the invention;

fig. 14 illustrates a mechanism for managing multiple instruction sets according to a preferred embodiment of the invention;

35 fig. 15 illustrates a tree of cascaded chains of service applications according to an embodiment of the invention;

fig. 16 shows the software components of a service support environment according to an embodiment of the invention;

5 fig. 17 shows steps performed by the service interaction module between the processing of the rule module and the processing of the service application in the embodiment of fig. 16;

10 fig. 18 illustrates the tree structure of the processing of rule modules according to an embodiment of the invention;

fig. 19 illustrates the recursive processing of rule modules in a situation where service applications generate new event contexts according to an embodiment of the invention; and

15 fig. 20 illustrates a mechanism of enforcing access control in connection with rule modules according to an embodiment of the invention.

20 Fig. 1 illustrates different types of feature interactions in a SIP service network. With the introduction of SIP, a new range of conversational and real time services is emerging on the Internet. These services 101-105 may be managed by end user terminals 106-108, also called user agents, or on one or multiple
25 intermediate network servers 109. Intermediate servers 109 may provide value-added services 102-103 to originating and/or terminating user agents, but also to the associated media client and server application(s). These intermediate servers can be proxy servers, redirect
30 servers, dedicated application servers or even user agents.

Communications messages sent between the user agents are routed by the intermediate server 109 via the Internet
35 110 or another communications network. A fundamental problem managed by an intermediate SIP server 109 is to

fulfil the user's service level expectations, to maintain good network performance and to allow flexible and scaleable definition of new services. Since performance is a key issue, the architecture should allow for some services to execute on the SIP server. Since it may not be feasible or possible to execute certain services on the SIP server, the architecture should allow for services to be executed on remote servers as well.

There is a need for flexibility of service definition, because the functionality of the SIP server should not be defined in advance. In particular, services may be continuously uploaded/registered from a range of sources including service providers and subscribers.

The owner of the SIP Server may, at the same time, be the SIP server provider, administrator and subscription provider. This stakeholder will be referred to as a network operator; it may be a kind of Telco or ISP. The network operator may own the domain name of the SIP server and provide service features to 3rd party service providers. The network operator installs service applications and rule modules on the SIP server, and offers services to subscribers. In this case the network operator acts as a service provider and service administrator. A subscriber has a contractual relationship with the network operator, in order to have a subscription, and to receive the subscriber services, offered by the SIP server. The subscriber may have a service that allows the subscriber to upload service applications and rule modules to the SIP Server. The subscriber owns these services for private use, i.e. personalized services. For simplicity all parties other than the network operator(s) and the subscribers will be referred to as 3rd party service providers which have a contractual relationship with the network operator.

Network server owners, network server providers, network server administrators, 3rd party service providers and subscribers are looking at the intermediate SIP servers as a potential platform for deploying services that cannot or should not be deployed in end-point user agents, for one reason or another.

There is a range of standardized protocols (e.g. SIP, SDP, SOAP), languages (e.g. CPL) and interfaces (e.g. SIP-CGI, 3GPP OSA API) that handle different aspects of service control. Also, services are likely to apply their control across multiple protocols, network components, languages and interfaces. The SIP server should be extensible to support all these aspects as required. These service applications may be owned by different parties, which have different authorization levels and different contractual relationships with the owner of the SIP server. The service applications may add value to the default processing of SIP requests and responses (but are not limited to doing so) at different specified call/session processing points, under different conditions, etc.

The following example illustrates that it is likely that multiple services may be invoked based on one SIP event. Consider an example where a subscriber, say Bob, has a SIP subscription with some SIP provider, say Telco. Bob has a range of services he would like to be placed in the network. Services like terminal independent services or services that become tailored to whatever terminal Bob is currently using. There may also be services that provide Bob with security, privacy and reliability. Also, Bob does not wish to manage his own services. Assume further that Bob works at a company, say Corp. The services that are invoked to manage an incoming INVITE message to Bob, may look like this:

The Telco may own and administrate the SIP server. The Telco also provides SIP services to subscribers, and hosts 3rd parties services. The Telco recognizes that Bob
5 is a subscriber on incoming INVITE. The first service is a Telco Call Barring application, to check if Bob has paid his last bill. The second service is a service provided by the Telco. It simply checks that the Caller's media codecs can be handled by Bob's current
10 location/terminal. If not the application will invite a media stream converter into the media stream flow (using 3rd party call control). This conversion is transparent for Bob, and the application will monitor for responses from Bob, and update the session descriptions as needed.
15 The third service is Bob's own callee preferences, e.g. a CPL script. This application monitors for responses for the proxied request, and possibly routes the INVITE to multiple destinations based on that. Say the INVITE is to Bob's current private SIP URL. On "no response" the
20 INVITE is proxied to his wife Alice. In some cases Bob wants all private calls to be diverted to his corporate SIP URL. The fourth service is also a Telco service, but provided to the Telco from a 3rd party application service provider, say besafe.com. This service checks for
25 message body content types and provides virus checking when needed, e.g. if there is an animated vCard included. The fifth service is an ISP service that offers the callee multimedia advertising in return of reimbursement of the Telco's charging. If the session is established
30 and, if it is a video- conference session, Bob receives a small streaming bar of information in the bottom of the video image. This service is a monitoring application and uses 3rd party call control. The ISP has an account on the Telco SIP server and can offer services directly to
35 the subscriber base owned by the Telco. Bob has subscribed to this service directly with the ISP, without

involving the Telco. The sixth service is a service managed by the company where Bob works. If the call is to Bob's current corporate SIP URL then the call is routed to Bob, based on data only known within the Corp's private LAN network. The last and seventh service is for administration purposes, as the Telco might want to do some logging.

The above example illustrates the diversity of contractual relationships associated with the different services triggered by an event. The services may be owned by different stakeholders, and implemented using different technologies.

Different services deployed in a service network may interact with each other. These interactions may be between services related to a single user, multiple users, or between customer services and system services. The interaction may further be between services running on the same network component or on different network component. In fig. 1, different types of interactions are illustrated: single user - multiple component (SUMC), customer - system (CUSY), multiple user - single component (MUSC), multiple user - multiple component (MUMC), and single user - multiple component (SUMC) interactions.

There are many causes of feature interaction, including "violation of feature privileges" and "violation of feature assumptions":

Violation of feature privileges: Privileges of features are broken, if features are invoked redundantly or by unauthorized parties. The main problem caused by violating feature privileges is redundant consumption of resources, or unauthorized access to resources. This may be un-

intended or malicious in nature. When violation of feature privileges occurs, there is a fight between features, authorized or non-authorized, for access to the resources. If features are invoked redundantly or by unauthorized parties, they may subsequently be the cause for violation of feature assumptions. Clearly, avoidance of violation of feature privileges is desirable.

As will be described below, according to the invention, violations of feature privileges are resolved by filtering on context, on contractual relationships, on conditions, on access control lists, privileges and rights.

Filtering on context relates to the need for interpreting an event in a certain context. This necessity becomes apparent when managing networked multimedia services that operate in a converged network.

Filtering on contractual relationships relates to the need for mapping an event to a set of features which are contractually obliged to process that event. This issue is particularly important in a de-regulated market where not only the provider of the network infrastructure and session control services can offer value added services to subscribers, but where 3rd parties have a legal right to make that offering as well.

Filtering on access control policies ensures that an event only causes authorized behaviour at the node.

Filtering on conditions ensures that a feature is not invoked redundantly when an event occurs and features to be invoked are detected based on the context, on contractual relationships and access privileges. These

conditions may depend on e.g. the properties of the event, system properties or network properties.

Violation of feature assumptions: Assumptions about
5 feature behaviour are broken, if the context of a feature is changed by another feature in such a way that the feature cannot work as intended. A violation of feature assumptions may cause ambiguous or conflicting behaviour.

10 Features are the visible behaviour of executing a service application. An instruction issued by the service application issues to control the value-added behaviour of the SIP node is termed a feature instruction. However, many features may apply their behaviour to a message
15 before the controlling instruction is sent back to the SIP node. The controlling instruction or instruction set that is sent back to the SIP node is termed a service control instruction. It contains a resulting event context, i.e. the properties of a SIP message that has to
20 be sent upstream or downstream in response to the original event context that triggered the service applications. Ambiguous behaviour occurs, when the service control instruction is different depending on the sequence in which the features gain control over the
25 current event context. Ambiguous instructions are not necessarily mutually exclusive.

As an example, let a SIP message body contain a color picture in gif format. Define S1 to be a service
30 application that is triggered on message body content of gif format. Let S1 be a service application that converts the gif format into a jpg format. Furthermore, define S2 to be a service application that also is triggered on message body content of gif format. Let S2 be a service
35 application that converts the gif picture into a black and white picture.

Let S1 be the first application to be invoked, triggered by the gif content. S1 will convert the gif picture to a jpg picture and write it to the current event context. S2 will never be invoked. The resulting event context will contain a color jpg picture.

Now let S2 be the first application to be invoked, triggered by the gif content. S2 will convert the color gif picture into a black and white picture and write it to the current event context. S1 will, subsequently, be invoked based on the gif content as specified by the current event context. S1 will convert the gif picture to a jpg picture and write it to the current event context. The resulting event context will contain a black and white jpg picture.

Clearly, S1 and S2 provide ambiguous behaviour, because the context of one feature is changed by the other.

Conflicting feature instructions are instructions that are mutually exclusive. Conflicting instructions will try to override each other. In this case conflicting instructions are typically also the cause of ambiguous behaviour.

Furthermore, all feature instructions are potential unauthorized instruction sets, unless they are explicitly authorized. Un-authorized feature instructions can have a malicious nature or be the result of a buggy service application. In any case they can do damage to the safety and integrity of the system and should be detected.

Monitoring service applications may cause additional problems to those already discussed. Monitoring service applications may issue asynchronous feature instructions

towards the SIP node at any time, as they are running continuously. As they are monitoring for future events, they may process on these events and provide more feature instructions. If there are multiple simultaneous
5 monitoring service applications, their generated feature instructions may depend on the order in which they are notified about the event. This adds to the complexity of the previously discussed problems.

10 The detection and resolution of violations of feature assumptions is much more complicated than the resolution of violations of feature privileges. As will be described below, according to the invention, means are provided to specify how to resolve ambiguous behaviour, means to
15 divide the feature interaction management in independent feature groups and administration domains, and a simple default rule to resolve feature interaction interference detected at run-time. Violations of feature assumptions are resolved by feature ordering based on the cascaded
20 chain principle and conditional triggering, and feature priorities based on the lock/unlock mechanism. This will be described in more detail below.

Other causes of feature interaction include, limitations
25 of network support, e.g. limited protocol functionality or limited user interfaces, intrinsic problems in distributed systems, such as resource contention, feature co-ordination, timing, or non-atomic operations, violation of system security, fraudulent, tampered or
30 eavesdropped messages, non-cooperative interactions by features with conflicting interests, etc.

In general three different types of solutions to feature interaction problems may be distinguished:

Infrastructure: The development of infrastructures for the deployment of features which integrate feature interaction management, i.e. which deal with the causes of some defined feature interactions. Feature interactions are managed both before (specification) and after (enforcement) feature deployment time, i.e. during specification or by enforcement of rule based policies, etc. According to the invention, the rule module scripts compose an infrastructure for deploying features which provides a framework for feature interaction management.

Service creation: The design of features with regard to causes of feature interactions, i.e. the detection of feature interactions during the design phase. This, feature interactions may be managed before feature deployment, e.g. via explicitness, feature interaction cause analysis, verification test, etc. According to the invention, requirements are imposed on the service creation and feature deployment strategy specification.

Run-time: The resolution of feature interactions as they occur. As not all feature interactions may be identified before feature deployment, they have to be detected at run-time, e.g. by cryptographic authentication, authorization and secrecy, rule based policies, resolution algorithms, AI negotiation, etc. According to the invention, simple rules are provided for how feature interactions detected at run-time are resolved. These rules include the checking of access control lists associated with rule modules, the resolution of access violations by alarm notification and taking violating rule modules out of operation, the checking of scripts specifying privileges and rights, and the resolution of violations of privileges and rights by alarm notification and taking violating features out of operation.

The general management process according to the invention includes the following steps:

1. Feature design specification, independently of other
5 features: A service is designed and specified without considering interactions with other features. However, according to an embodiment of the invention, the service is designed with respect to the processing points model described below.
- 10 2. Contract negotiation: The party wishing to deploy a service in the service network negotiates with the administrator of the service network which privileges and rights may be granted to the service. According to an
15 embodiment of the invention, a privileges and rights script associated to the service is created.
3. Feature interaction analysis: The possible feature interactions are investigated based on experience and possibly knowledge about some of the features deployed in the service network.
- 20 4. Feature deployment strategy specification: The author of a rule module which is to deploy the service acquires an administrative domain. Based on the feature interaction analysis and knowledge about the service, subscribers and users, the feature deployment strategy is
25 specified in a rule module script.
5. Feature implementation.
6. Verification test.
7. Feature installation and activation.
8. Feature run-time behaviour and management.
- 30 9. Feature de-activation and de-installation.

According to the invention, a framework for feature interaction management is provided which

- allows for the analysis stage to be easily mapped to
35 specification rules by providing a simple language and framework with easily understood principles,

- defines clear boundaries between the groupings of features which are known to a given analysis entity and those which are not,
- provides simple and thus easily understandable rules for the handover of control between these groups of features. When handing over from one group of features to another there is a mechanism for ensuring that a subsequent group of features does not compromise the previous group, and
- simplifies the analysis stage by specifying processing points in the processing of events at which there are guaranteed pre-conditions and at which applications are only allowed to give certain instructions to the server.

15

Fig. 2 illustrates the network elements involved in a service environment SIP server architecture according to an embodiment of the invention. The previous sip client 203 represents any client, such as a SIP enabled PC, a wireless terminal, a previous hop proxy, a SIP/PSTN gateway etc. The client makes requests for session services with incoming requests to the SIP server 202. The SIP server 202 represents the proxy, redirect or dedicated SIP enabled application server where the service support environment 201 is implemented. Alternatively, it may be any other SIP enabled entity that triggers value added services, such as a user agent, registrar, or the like. Services located in the SIP server service support environment 201 are defined by service applications, such as SIP-CGI scripts, rule modules and service features. The SIP Node 202 may hand over control to the service support environment 201 on reception of an event. The service support environment 201 can subsequently invoke a relevant service application according to certain filter criteria and based on that event. The service application returns a

feature instruction or a set of instructions. The service support environment 201 hands back control to the SIP node 202, together with service control instructions that informs the SIP node 202 about how to process the event.

5 The SIP node forwards the request to the next SIP client 204. Responses to the request will be routed in the opposite direction from the next SIP client 204 via the SIP node 202 to the previous SIP client 203, possibly triggering additional services. The next sip server 204

10 represents any server, e.g. a SIP enabled PC, a wireless terminal, a next hop proxy, a SIP/PSTN gateway etc. The server handles incoming requests. The remote server 206 offers remote service execution, e.g. in another service support environment. Based on e.g. performance criteria

15 the service support environment 201 may initiate processing on the remote server 206, e.g. by use of request/response protocols. In this way different categories of services may be invoked and managed on different hosts. The protocol used towards the remote

20 server may be any protocol supporting request/response dialogs, e.g. SIP, ICAP, HTTP, OSA API, etc. The administration server 205 performs administrative tasks on the service environment. It is responsible for configuring the service support environment 201, which is

25 associated to the domain of the SIP Node 202. Hence the environment of the service support environment 201 includes a SIP node, service applications, remote hosts and an administration entity.

30 Fig. 3 schematically illustrates a system architecture of a SIP server which supports a service execution rule mechanism according to an embodiment of the invention. The service support environment 301 enforces the deployment infrastructure, i.e. how one feature interacts

35 with another. Hence, the service support environment 301 provides the functionality in the SIP server that

supports value-added services. The service support environment 301 comprises a rule engine 303 for managing the rule modules 308 stored in an execution rule base 307. The rule engine 303 further comprises a rule module
5 execution module 314 responsible for processing rule modules. The rule engine can invoke services 309-311 via the service execution engine manager 302. The service definition manager 312 provides functionality for the administration of the rule modules. The SIP server
10 further implements a SIP protocol stack 304 including the SIP default functionality 306 and a SIP message parser 305. The message parser 305 supports the SIP protocol and extracts message properties that can be interpreted by the rule engine 303. The administration server 313
15 performs administrative tasks on the service environment.

An important mechanism for implementing service deployment policies according to the invention is the specification of deployment rules as service execution
20 rules specified in rule modules. Service execution rules specify conditions and actions that need to be taken, if the conditions are fulfilled. According to the invention, a programming language for specifying these rules is defined which will be referred to as Service Execution
25 Rule Language (SERL).

The rule modules 308 are managed and executed by the rule engine 303. The rule engine 303 is the main functional entity that implements the triggering and feature
30 interaction mechanism, and is part of the service support environment 301. When an event occurs, the message parser invokes the rule engine 303 and hands over the event to the rule engine 303. The rule engine 303 finds and loads the relevant rule modules 308 and processes those that
35 are relevant to the received event in the correct order. The filtering includes a detection of contractual

obligations. The events define the context in which rule modules are processed, i.e. the conditions are evaluated according to the properties of the SIP message events. The rule engine 303 invokes the corresponding actions

5 when the rule pattern matches the given message properties. Based on the content of these actions, the rule engine 303 may issue invocation instructions to the application execution engine manager 302 or another appropriate entity within the service support system.

10 SERL scripts do not have knowledge about the services that they invoke and manage, other than the knowledge about how to invoke them and manage features in general. The controlling instructions received from the invoked service applications are mediated back to the SIP stack

15 304 when the last rule module has been processed.

The rule engine 303 further manages the information which is sent between different services. When an event occurs, an event context is established containing the relevant

20 properties of the event in a standardized way. The message properties have a name and a value. The value may be determined by the message. Examples of SIP message properties are

25 Name: sipRequest.Request-URI, Value: sip:bob@corp.com
Name: sipRequest.To, Value: Bob Smith <sip:bob@corp.com>
Name: sipResponse.Status-Code, Value: 301

An example of a SDP message property is:

30 Name: sdp.m, Value: video 48232 RTP/AVP 0

The rule engine 303 supports a number of internal APIs that can be accessed by the interfacing functional

35 entities. These APIs include

- a message notification API used by the SIP stack 304 for message notification from the message parser 305 to the rule engine 303,
- a rule base definition API, used by the service definition manager 312,
- a service instruction API used by the application execution engine manager 302 to hand over instructions to the rule engine 303 on behalf of the service applications 309-311, and
- an arming API used by the application execution engine manager 302 to request the arming of triggers and transaction events on behalf of the service applications 309-311.

The application execution engine manager 302 embeds and manages a number of application execution engines 315 for different types of service applications, e.g. a OSA engine, a CPL engine, CPL interpreter, a CGI engine and/or a Servlet engine. It supports the interface induced by the rule engine 303 and maps between the application API and the rule engine API. From the viewpoint of the rule engine 303 all application execution engines 315 look like a single entity, i.e. the application execution engines manager 302. The service definition manager 312 may provide further functionality to the application execution engines 315.

The APIs provided by the application execution engines manager 302 include:

- an invocation API used by the rule engine 303 for invocation instructions from the rule engine to the application execution engine manager 302, and
- a notification API used by the rule engine 303 for notification instructions from the rule engine to the application execution engine manager 302.

The default SIP server behaviour 306 comprises the functionality of a proxy server, redirect server, application server or even a user agent. It may further include Registrar or IM&P functions. It provides an
5 interface which can be accessed by the rule engine 303 to place instructions. Alternatively, a service application may implement the SIP server behaviour. Therefore, it is possible for the rule engine to invoke only parts of each of the above functions as required by the service
10 applications.

The service definition manager 312 provides an O&M API used by the administration server 313, service applications 309-311, and the SIP stack 304. The API
15 provides functionality for

- Manual authentication and authorization of new rule modules and service applications.
- Manual configuring of rights and privileges of rule modules and service applications.
- 20 - Manual loading of rule modules and service applications.
- Manual activation/deactivation of rule modules and service applications.
- Manual validation of rule modules.
- 25 - Manual validation of service applications implemented using script languages.
- Manual deletion of rule modules and service applications
- Manual listing of all rule modules and service
30 applications together with their status.
- Manual modification of rule modules.

The service definition manager 312 may further provide an interface supporting an automatic handling of some of the
35 above manual functions and/or additional features, such as listing available service features, getting interface

and/or version of a service feature, listing supported execution engines and script languages, such as JVM, CPL, Perl, etc., installation/un-installation of service features, activation/deactivation of service features, 5 registering of service application that will provide services as service features to other service application, subscription to service features, invoking/terminating of service features, statistic operations like "monitor processor load", "busy hour 10 calls", accounting operations, activation/deactivation, reading, resetting of accounting records, etc.

Here, the term service features comprises functions which are offered to service applications. The service features 15 are considered as being integrated into the service support environment 301 and SIP stack 304. The application execution engine manager 302 and the service definition manager 312 both provide service features to the service applications 309-311. Although some of the 20 features offered by the application execution engine manager 302 are mediated through the rule engine 303 and the SIP stack 304.

The service applications 309-311 are programs, compiled 25 or interpreted, executing in the service support environment 301 of the SIP server, or on a remote server. Their purpose may be related or unrelated to the basic functions of the SIP sever. Service applications may implement SIP server behaviour. Service applications may 30 offer services to other service applications, i.e. they may be service features. For example, a service application may contain some MIME type converter, and offer it as a service feature to other service applications. Other service applications may then use 35 this function in order to perform a service. Preferably, service applications should be portable between SIP

servers and may be executed on remote servers. Service applications may access, or may be accessed through global or local naming convention, a standardized or local namespace, by using specified file paths to the applications, etc., open and standardized API, e.g. SIP-CGI, OSA API, HTTP, ICAP, CPL, servlets, etc.

A standardised mechanism of accessing service application is an advantage for 3rd party service providers.

10

The SIP Server does not necessarily manage remotely placed service applications, as it may have no knowledge about them. In this case, the SIP server or the application execution engine manager 302 may require location information provided by the triggering information.

The message parser 305 is responsible for interpreting messages received, isolating well-defined elements as message properties, and causing actions to be activated when appropriate. When an event occurs, a message object is established containing the relevant properties isolated by the message parser 305.

25 Preferably, any supported protocol has a corresponding message parser. A parser may contain subordinate parsers that correspond to subordinate protocols, i.e. embedded in other protocols like SIP. For example, within a SIP message parser there may be separate parsers for handling different media types such as SDP, HTML, XML, and XHTML. From the standpoint of the rule engine, all parsers look like a single engine. The API for message parser should preferably be defined such that parsers for new protocols and content can be added modularly.

35

For any protocol to be supported by the service support environment 301, the interface between the message parser of that protocol and the rule engine covers

- the set of properties defined by the message parser, including the property name, its relationship to the message it characterizes, and the ability of an action to modify it, and
- the processing points at which rules can be activated.

10 In one embodiment of the invention, the service support environment 301 or each of the functional entities within the service support environment may be placed at separate hosts serving multiple servers, e.g SIP, Web, WAP, I-Mode, RTSP servers, and accessing multiple application
15 servers, e.g. databases, OSA, Web, SIP, etc.

For example the rule engine 303, the application execution engine manager 302 and the service definition manager 312 may be located at each their host, possibly
20 interfacing via IP. The interfaces between the rule engine and the various servers may be standardized or proprietary. The interfaces between the rule engine, the application execution engine manager and the service definition manager should, preferably, be proprietary and
25 packaged together. The interfaces between the service application servers and the service support environment are preferably standardized, like OSA API, HTTP, SIP or some database API like some SQL based API. In such a distributed configuration the a further advantage of the
30 invention becomes apparent, as the rule engine is able to manage many types of events, not only SIP events, but also HTTP events, SMTP events, Wap events, Media codec events like MPEG7 events or 3D virtual reality or gaming object events, etc.

Still referring to fig. 3, the trigger mechanism according to an embodiment of the invention may be illustrated by a simple example where it is assumed that only a single rule module is relevant for a received event. In figure 3, the numerals in circles refer to the following steps of a trigger example:

1. A SIP request is received at the message parser 305.
2. The SIP message parser 305 generates a SIP message object. The rule engine 303 converts this into a SIP event context.
3. The relevant rule modules 308 are located in the execution rule base 307 based on the event context, processing points and contractual relationships. An embodiment of this mechanism will be described below.
4. The rule engine 303 loads the found rule module 308 from the execution rule base 307 into the rule module execution engine 314, and executes it. As an example, let the loaded rule module contain the following functionality, i.e. a rule specifying trigger criteria and invocation actions for service applications 310 and 311:

```
<rulemodule priority="1" >  
  <ancillary>  
    <owner class="Network Operator">  
      <name>Telco</name>  
      <id>telco.com</id>  
      <id>123.123.123.123</id>  
    </owner>  
  </context>  
  <protocol>  
    <SIP version="2.0" />  
  </protocol>  
</context>  
</ancillary>
```

```
<rule processing-point="1">  
  <property name="request-line" matches="INVITE^">  
    <property name="TARGET" matches="^telco.dk">  
      <action> "invoke application Y" </action>  
5      <action> "invoke application Z" </action>  
    </property>  
  </rule>  
</rulemodule>
```

- 10 5. A trigger is reached, and the specified service application Y (310) is invoked by sending an invocation command to the application execution engine manager 302.
- 15 6. The application execution engine manager 302 locates the application Y (310).
7. The application execution engine manager 302 loads the relevant service application Y (310) into an application execution engine 315 and, subsequently, activates it for execution.
- 20 8. The resulting instruction from application 310 is handed back to the rule engine 303.
9. The rule engine 303 resumes processing the rule module, and triggers another application Z (311). An invocation command is given to the application execution engine manager 302.
- 25 10. The application execution engine manager 302 locates the application Z (311).
11. The application execution engine manager 302 loads the relevant service application Z (311) into an execution engine 315 that can run the application and, subsequently, activates it.
- 30 12. The resulting instruction from application 311 is handed back to the rule engine 303.
13. The rule engine 303 resumes processing the rule module, and finds out that the rule module has no more rules to execute, and that there are no more
- 35

rule modules to load. Subsequently, the rule engine 303 sends the final result of the services to the SIP default behaviour 306.

14. The SIP default behaviour 306 merges the output from the rule engine 303 with possible default output and sends the SIP message to the message parser/converter 305.

15. The SIP Message is e.g. proxied.

10

In the case of multiple rule modules, the above steps 13-14 may alternatively be as follows

13. The rule engine 303 resumes processing the rule module, and finds out that the rule module has no more rules to execute. The rule engine 303 searches for rule modules with lower order priority than the previous executed rule module.

14. Go to point 3 in the previous example. If a rule module is found it will be run as in previous example, possibly invoking other applications, e.g. the application 309 will be invoked.

When no rule modules or services are installed to control a session, the rule engine 303 hands over control to the SIP server, and specifies an empty output. The SIP server may possibly merge the empty output with the default behaviour 306 of the server as in SIP-CGI, according to default SIP behaviour for Registrars, Redirect Server and Proxy Server.

Fig. 4 illustrates the structure of a rule module according to an embodiment of the invention. According to the invention, a rule module 401 is conceptually a tree comprising a number of service execution rules 401-402, each rule specifying a number of conditions 403-404 and

405, respectively, and a number of actions 406-407 and 408, respectively, that need to be taken if the corresponding conditions are fulfilled. In the following, conditions will also be referred to as patterns. At a high level a rule module further comprises

- owner information 409 specifying the owner of the rule module, i.e. an identifiable party with an interest in the SIP node. This party may own multiple rule modules. The owner information may further include information about contractual relationships.
- protocol information 410 specifying the context in which to interpret the rules in the rule module. Examples of protocols are SIP, SDP, HTTP, H.323, etc. Preferably, there should only be a single protocol defined per rule module. Otherwise some overlap may provide ambiguous interpretation. For example, both SIP and HTTP have content type header fields.
- access control information 411 specifying which parties have the right to invoke, administer and read the rule module,
- a rule module identifier 412, preferably a unique identifier. Additionally, a rule module may include a number of aliases.

A rule module may further comprise ancillary information 413 providing further context information for a rule module and index information 414.

A pattern 403-405 is an expression that can be evaluated with respect to the message properties in an event context, and either the rules will match or fail to match the properties in the context. Actions 506-408 may identify applications, built-in service features, remote servers, load sharing hosts, or next hop sip servers to be invoked. They may further identify downloaded and externally placed action objects, etc. Rules, conditions

and actions may have parameters that describe their behaviour, they are the attributes of the tree nodes. The branches have a weighting indicating which branch should be processed first.

5

Preferably, every rule module has an explicit ordering priority assigned to it, i.e. the order priority specifies the sequence in which rule modules are loaded by the rule engine.

10

According to a preferred embodiment of the invention, the nodes of the tree in the graphical representation are represented by XML elements. The branching from one node to the next nodes is represented by enclosing the element representing the next node(s) within the element representing the node branched from. The weighting of the branches is given by the order in which they are represented in the script. This weighting gives the order in which to process the elements when an event is received. In other words, the script is executed from top to bottom. Thus, in this embodiment, there are no loops within a SERL script. Preferably, the format of a rule module is specified as an XML DTD or XML schema.

25 It is an advantage of this embodiment, that it is based on a standardized and extensible language to specify languages.

30 In one embodiment of the invention, policies can be associated with rules specifying privileges and rights. Each node of the rule module may have an associated Policy node, which may contain access control lists.

35 The following is an example of an instance of a rule module specified as a SERL script:

```
<?xml version="1.0" ?>
  <!DOCTYPE serl PUBLIC "-//IETF//DTD RFCxxxx SERL
    1.0//EN" "serl.dtd">
  <rulemodule priority="4">
5    <owner class="service provider">
      <name>Third Party Example</name>
      <hostname>sip.example.com</hostname>
      <company>www.3party.example.com</company>
    </owner>
10   <protocol protocolname="sip" protocolversion="2.0"
      />
      <rmid>4444</rmid>
      <serule processing-point="1">
        <property name="sipEventContext.TARGET"
15        matches="From">
          <property name="sipRequest.method"
            matches="INVITE">
            <property name="User-Agent" matches="company1">
              <action>
20                <invoke type="soap">
                  <objname>company2</objname>
                  <objoperation>Event_Report</objoperation>
                  <objaddr>
                    www.company2.com/soap/servlet/rpcrouter
25                </objaddr>
              </invoke>
            </action>
          </property>
        </property>
30    </property>
      </serule>
    </rulemodule>.
```

Fig. 5 illustrates the grouping of services into
35 constrained sets of services according to an embodiment
of the invention. According to this embodiment, services

501-508 are grouped into a set of feature groups 509-510. In the example shown in fig. 5, feature group 509 contains features 501-504 and feature group 510 contains features 505-508. For each feature group, certain
5 constraints are imposed on the services of that group, i.e. they are only allowed to issued well-defined types of instructions. The feature groups do not specify how single features are allowed to interact within a feature group, as long as they do not violate the constrains on
10 the group behaviour. The feature groups are sequentially ordered, e.g. by enumerating the feature groups. In the example of fig. 5, feature group 509 is the K-th feature group in a sequence of feature groups and feature group 510 is enumerated by K+1. The ordering of feature groups
15 imposes an order of execution, such that the features of feature group 509 are executed prior to the execution of the features in group 510. Preferably, the definition of a feature group comprises a specification of the group ordering, e.g. by specifying a group index, by specifying
20 a previous and a next feature group, or the like. Preferably, the definition of a feature group further comprises a specification of pre-conditions, i.e. conditions on which the features of that feature group may rely, and a specification of constraints enforced on
25 the behaviour of the services of that feature group. Hence, the feature groups are an ordering mechanism formalizing a fundamental type of feature interaction by providing an ordered sequence in which service applications are invoked. Therefore, this mechanism
30 provides a framework for grouping the problem of feature interaction into constrained sets of features. The mechanism actually solves feature interactions between these feature groups, because, due to the constraints on the feature groups, the feature assumptions when handing
35 over control from one group to the next, is deterministic and well-defined. It is an advantage of this grouping

that it provides a mechanism for decomposing the problem of feature interaction analysis into smaller, independent problems, thereby making the problem easier to manage. Furthermore, it is an advantage that formalized problem areas may be delegated to independent parties, making the problem more scalable. The mechanism of feature groups gives the administrator and service provider the ability to categorize the service applications to different points in processing time, where the service applications should be invoked.

Fig. 6 illustrates the grouping of services into constrained sets of services corresponding to locations in the round trip SIP message flow according to a preferred embodiment of the invention. According to this embodiment, the feature groups are related to locations P1-P6 in the logical request/response round trip message flow which have certain pre-conditions guaranteed for the event context, and where a service application can be invoked based on the constraints about which behaviour is allowed at that location. These feature groups will be referred to as processing points.

The six processing points P1-P6 are a general grouping of services. Processing points P1-P3 and P4-P6 logically cover corresponding problems. Processing points P1-P3 include services which are triggered on requests 607 and processing points P4-P6 include services which are triggered on responses 608. Logically, processing points P1 and P4, points P2 and P5, and points P3 and P6 correspond to corresponding constraints.

A SIP message can roughly be divided into two parts, the signalling properties, i.e. properties related to SIP, SDP, etc. and the message body content which is not related to signalling, such as gif, html, etc. Services

invoked on a SIP event can therefore be grouped into those which:

- impact the signalling properties,
- 5 - impact the non-signalling message body content only, and
- those which neither impact the signalling properties nor the non-signalling message body content.

10 This gives the basic grouping of services into three groups 601-603 associated to the three processing points P1-P3, respectively. Similarly, in the response pass, the corresponding groups 604-606 are related to the processing points P4-P6, respectively.

15 Here, the term signalling properties comprises SIP and SDP message properties that can be matched in a rule module condition to invoke a service. Consequently, applications that change a signalling property may cause
20 another application to be invoked, which might invoke yet another application and so fourth. The advantage in moving the handling of service applications which do not change signalling properties to a processing point after all changes to the signalling properties has occurred, is
25 that the administrator has an easier task in ordering the applications at the processing points P2 and P5.

In one embodiment of the invention, the processing points P1-P6 may be defined as follows:

- 30 - Processing point P1 - Previous Hop Client Request: A SIP request 607 from a previous hop client has been received. For this particular request and for this particular subscriber no rule modules have been invoked at any previous processing point. This
35 processing point comprises services which impact the

SIP/SDP signalling properties of the resulting messages(s), but not the message body content.

5 - Processing point P2 - Request Content Point: The
signalling properties of the resulting SIP/SDP
message(s) have been generated and cached, i.e. this
part of the SIP message is ready to be sent.
Additional updates to the call/media control
signalling properties should not occur, unless
explicitly authorized. The services 602 invoked at
10 processing point P2 should not impact the SIP/SDP
signalling properties as generated in processing point
P1. However, there may be exceptions to this rule. For
example, services may impact SIP headers like Alert-
Info, Call-Info, Content-Disposition, Content-
15 Encoding, Content-Language, Content-Length, Content-
Type, and Require. Special care should be taken, if
these signalling properties are changed at both
processing points P1 and P2. Examples of types of
media content that may be processed at this processing
20 point include SOAP, HTML, vXML, SMIL, gif, mpeg7, au,
etc.

- Processing point P3 - Request Batch Point: The SIP
resulting message(s) has been generated and sent. No
more updates to the resulting message(s) can occur.
25 This processing point corresponds to services 603 that
just need to be invoked but do not produce and output
which is needed to process the request. These services
may only read the resulting message(s), not update it.

30 The processing points

- P4 - Previous Hop Server Response,
- P5 - Response Content Point, and
- P6 - Response Batch Point
may be defined for response messages 608 analogously to
35 the processing points P1-P3, respectively.

Alternatively or additionally, other processing points may be defined. For example, an additional processing point P0 (not shown) may be defined without constraints on the features and which is triggered on any event, i.e. on the reception of both requests 607 and responses 608.

Preferably, additional processing points may be associated with one of the high level processing points. For example, sub-processing points may be defined for processing point P1, which may be named P1.1, P1.2, P1.2.1, P1.2.2 etc. Preferably, new processing points should only be defined, if they define a logical group of services that may be beneficial to invoke according to specified pre-conditions.

The following table includes some examples of services which may be defined at the different processing points at a originating and terminating SIP server, respectively:

PP	services on originating server	services on terminating server
P1	<ul style="list-style-type: none">- Originating call barring- Caller preferences/CPL call barring- MM adviser controller (reimbursement)- Privacy/secrecy enforcement- Anonymous caller	<ul style="list-style-type: none">- Terminating call barring- Callee preferences/CPL routing- Outlook based routing decisions- Presence based routing decisions- Automatic instant messaging services- Media stream converter controller- Online gaming server
P2	<ul style="list-style-type: none">- Animated Vcard- Greetings of the day	<ul style="list-style-type: none">- Virus scanner

P3	- Logging	- Logging
P4	- Call back service if queuing	- IVR controller - Music queue controller - MM advertiser controller - privacy/secrecy enforcement
P5	- Vcard converter - Virus scanner	- Animated Vcard - Joke of the day
P6	- Logging	- Logging

Fig. 7 illustrates the processing flow between services belonging to different processing points according to an embodiment of the invention. In the embodiment described in connection with fig. 6, services that impact different parts of a SIP message are grouped into different processing points. This may be utilised in order to provide an efficient processing of services. Services which impact the SIP signalling properties of an incoming message 701 are included in processing point P1. Therefore, the output from processing point P1 comprises the final call/media control instructions 702. This can immediately be merged with possible SIP server default behaviour and handed over to the message converter 705 which prepares the outgoing SIP message 703. In addition, the services invoked at processing point P2 can rely on that the signalling properties can no longer change, e.g. they can with confidence apply services that dependent on the final generated destination address.

By concentrating services that impact the non-signalling SIP message body content only (and possibly the related content header fields) in processing point P2, the output 704 from processing point P2 is sufficient for the server to actually send the resulting outgoing SIP message(s) 703. The output 704 from processing point P2 is handed over to the message converter 705, merged with the output

702 from processing point P1, and the message 703 is sent.

When the output from processing points P1 and P2 is
5 ready, processing point P3 is reached. The services in
processing point P3 do not impact the content of the
resulting SIP message 703, but may rely on it. Hence,
preferably, the resulting SIP message generated from
point P1 and P2 may be sent before waiting for the
10 execution of processing point P3 services, thereby
increasing the efficiency of the system.

Hence, the grouping of services according to processing
points has the following advantages: Processing points
15 simplify the problem of feature interaction. Processing
points make the problem of feature interaction more
scalable. Processing points improve latency of processing
messages when using parallel processing, i.e. multiple
processors. Responses or proxy requests can be handed
20 back to the network faster, because service applications
that do not specify instructions to the SIP Server are
located in processing point P3. Hence, the response or
proxy request does not need to wait for the service
applications in processing point P3 to be invoked.

25 Fig. 8 illustrates the grouping of services into rule
modules corresponding to administrative authority
according to an embodiment of the invention. As rule
modules have a rule module owner associated with it, they
30 correspond to an administrative domain where an
administrator, i.e. the rule module owner, can specify
service deployment policies. Within a rule module, the
order of actions is assigned by the owner of the rule
module, i.e. the SERL script author.

According to this embodiment of the invention, each rule module has a priority assigned to it. The priority specifies the relationship between rule modules. The higher the rule module priority, the earlier the rules in the rule module are applied. Fig. 8 schematically shows two rule modules 801 and 802. Rule module 801 is owned by administrator A, while rule module 802 is owned by administrator B. The rule module priorities of the rule modules 801-802 are indicated by a rule module order, where a low order corresponds to a high priority and vice versa. For example, rule order 1 may be defined to be the highest priority. In the example of fig. 8, rule module 801 has a rule module order h and rule module 802 has an order h+1. Consequently the services 803-804 invoked by rule module 801 are executed before the service 805 of rule module 802.

Rule modules with the same owner may have the same ordering priority. Preferably, in that case, they should have different event contexts (e.g. SIP, HTTP, ...), thereby avoiding that more than one rule module with the same priority is invoked on the same event context. In other words rule modules that can be invoked on the same message property from a given event context should have different priorities.

Each administrative domain is independent of other administrative domains. This means that one administrative domain does not need to have knowledge of the services deployed in other domains. Each administrator is responsible for analyzing and specifying service deployment policies in one given domain.

If more than one rule module or service application is invoked at a given event, they each will provide a set of feature instructions that will specify how this event is

to be processed, for example whether the event should be forwarded or terminated. Clearly such feature instructions should not be processed simultaneously, and some feature instructions may override others. However,
5 in some cases it might not be important in which order feature instructions are applied or whether they are applied simultaneously.

According to the invention, a mechanism is provided to
10 allow administrative domains to protect their feature instructions when handing over control to another administrative domain. This means that each administrative domain may limit the hand-over of control to properties that are unimportant to the correct
15 behaviour of its features. The object that is passed from one service to another is the current event context 806-807. Preferably, when handing over the control of the event context 806 between service applications belonging to the same administrative domain, the default rule is
20 that no properties of the event context is protected. If something has to be protected, it has to be done explicitly. Between administrative domains, however, all properties of an event context 807 are protected by default. If something is not necessary to protect between
25 administrative domains, it has to be explicitly marked un-protected. Preferably, the right to mark properties protected and/or un-protected should be governed by a privilege associated with a rule module.

30 In other words, higher priority rule modules may lock message properties if they have the privileges to do so, and locked message properties cannot be unlocked by rule modules unless they have the privileges to do so.

The following example of a fragment of a rule module illustrates the locking/unlocking of message properties after a feature F1 has been executed:

```
5  <rulemodule priority="1">
    ...
    <serule processing-point="1">
        <action>
            <invoke type="sip-servlet">
10         <objname>
                F1
            </objname>
            <lock object="property"
                name="sipRequest.Request-URI" />
15         <unlock object="keep" />
            </invoke>
        </action>
    </serule>
</rulemodule>.
```

20 In one embodiment of the invention, an action output may terminate the rule module, i.e. the downstream applications are not invoked. Furthermore, actions may explicitly set privileges on message properties in the

25 current event context, if they have the privileges to do so. It is an advantage of this embodiment that the administrative domains are independent once they have been assigned a priority.

30 According to this embodiment, an overall administrator assigns the priorities or ranges of priorities of the rule modules based on contractual relationships with each of the rule module owners. The overall domain administrator would naturally also be the owner of the

35 highest priority rule module. For example, at the top of the hierarchy is the SIP service provider who owns the

domain name and IP address of the host. The SIP service provider may provide many services to many parties. He may also run applications for his own purposes, for example logging, accounting, statistics gathering, fraud
5 detection, advertising etc. He may place one or more rule modules on the server.

For example, the SIP service provider may place a rule module for each of his subscribers. When a SIP event was
10 received in the SIP server the SIP service providers main rule module may invoke some of his own services and then hand over control to the rule module for the appropriate subscriber. In this rule module there may be rules to trigger services tailored for the specific subscriber.
15 These services may be provided by the SIP service provider himself or bought-in by the SIP service provider from third party service providers. The SIP service provider would be responsible for analyzing feature interaction problems between these services and
20 specifying order priorities and instruction priorities between the service applications, preferably by using SERL. In this case the main rule module and the subscriber rule module are both owned and administered by the SIP service provider. They are said to be in one
25 administration domain. At some point in the order priority the SIP service provider may decide to hand over to a rule module owned by the subscriber himself. Note that individual subscribers may be able to write their own SERL scripts, or they may only able to update
30 preferences, e.g. via a HTML form, from which a SERL script may be generated. The subscriber's rule module may invoke a CPL script, or may in turn invoke another rule module, some third party service provider's rule module, or the like. If the subscriber's rule module includes
35 more than one action, the subscriber has to specify which action should be done first. Likewise in the third

parties rule module the actions should be ordered according to the wishes of the third party.

5 In another scenario, the subscriber may be an employee of a corporate customer. In this case the SIP service providers rule module may first hand over to the corporate customer's rule module which might invoke some applications and then hand over to the individual subscriber's rule module.

10

The top-level administrator may like to intersperse the priority order with a number of rule modules. However, it may be cumbersome to intersperse rule modules with lower priority rule modules of other parties. One solution to
15 this problem is to separate rule module priorities with sufficient back-up ranges that can be used at a later time. Alternatively, a hierarchical administration domain model may be applied, as will be described in connection with fig. 12.

20

Preferably, the service support environment supports a mechanism for notifying administrators when an application attempts to alter a protected property. Preferably, an application that attempts to alter a
25 protected property is taken out of service.

It is noted that modifications of the above scheme are possible. For example, in a hierarchy of administrative domains each level administrator may have complete scope
30 to arrange rule modules at levels with lower priority, within an allocated priority range. The higher level administrator would delegate this scope.

It is an advantage of this embodiment that it provides a
35 scalable mechanism for managing and implementing service deployment policies. According to this embodiment, the

service support environment is able to host 3rd party service deployment policies without considerable extra work for the domain administrator. It is a further advantage that limitations are placed on the number of stakeholders or the number of contractual relationships between them.

It is a further advantage of this embodiment that the task of feature interaction analysis is broken down into administrative domains, thereby allowing a distribution of the problem between the parties concerned.

Fig. 9 illustrates an embodiment of the invention where services are grouped both according to processing points and according to administrative authority. Fig. 9 schematically shows a rule module 901 with priority order h and a rule module 902 with priority order $h+1$. Rule module 901 invokes services $F1$, $F2$, $F5$, and $F6$, while rule module 902 invokes services $F3$, $F4$, $F7$, and $F8$. The services $F1$ - $F8$ are further related to processing points 903 and 904. Processing point 903 includes services $F1$ - $F4$, and processing point 904 includes services $F5$ - $F8$. The processing points 903-904 are enumerated such that processing point 903 has the index K and processing point 904 has the index $K+1$, indicating that processing point 903 is processed before processing point 904.

According to this embodiment, rules are applied in sequence of the processing points. Consequently, when processing rule modules, only the service execution rules belonging to the current processing-point are executed.

For each processing point, there are sets of rule modules which are grouped according to priority and which may be invoked at that processing point. All priority 1 rule modules are grouped in set 1, all priority 2 rule modules

are grouped in set 2, and so on. However, given an event context, at most one rule module is invoked in each of these sets. A rule module may be distributed over multiple processing points.

5

In the example of fig. 9, first the services of the processing point 903 are processed according to their rule module priority, i.e. services F1 and F2 are processed before services F3 and F4. Subsequently, the features F5 and F6 of processing point 904 and priority h
10 are invoked before services F7 and F8.

The rule modules 901 and 902 which are schematically illustrated in fig. 9, may be expressed as indicated in
15 the following examples of rule modules.

Rule module 901:

```
<rulemodule priority="h">
  <owner> <name> A </name> </owner> ...
20  <serule processing-point="k">
    <property ...>
      <action> F1 </action>
      <action> F2 </action>
    </property>
25  </serule>
    <serule processing-point="k+1">
      <property ...>
        <action> F5 </action>
        <action> F6 </action>
30    </property>
      </serule>
    </rulemodule>.
```

Rule module 902:

```
35 <rulemodule priority="h+1">
  <owner> <name> B </name> </owner> ...
```

```

    <serule processing-point="k">
      <property ...>
        <action> F3 </action>
        <action> F4 </action>
5      </property>
    </serule>
    <serule processing-point="k+1">
      <property ...>
        <action> F7 </action>
10      <action> F8 </action>
      </property>
    </serule>
  </rulemodule>.
```

15 Fig. 10 illustrates the processing mechanism of the embodiment of fig. 9 in the case of multiple rule modules and multiple processing points. According to this embodiment, there are four processing points enumerated P0 through P3 and five rule modules RM A through RM E, 20 with priorities 1 through 5, respectively. According to the processing rules described in connection with fig. 9, services are processed as indicated by the circles in fig. 10 which are enumerated 1 through 20 and connected by arrows. Each enumerated circle may represent a set of 25 rules invoking a number of actions.

Fig. 11 illustrates another example of the processing rules described in connection with fig. 9. Fig. 11 illustrates instances 1111a-1115a of rule modules RM A 30 through RM E on the left side 1101 and instances 1111b-1115b of the rule modules RM A through RM E on the right side 1102 of the figure, respectively. The rule modules RM A through RM E comprise rules corresponding to processing points P0 through P6. Hence, the instances of 35 rule modules on the left side 1101 may represent different sections of the same respective rule modules as

the rule modules on the right side 1102. As was described in connection with fig. 6, the processing points P1-P3 are triggered by SIP requests, while the processing points P4-P6 are triggered by responses.

5 Processing point P0 is triggered by both requests and responses. Hence, in this example, an incoming SIP request 1103, e.g. an INVITE request, triggers services of rule modules RM A through RM E which are related to processing points P0-P3, in the order indicated by the

10 arrows on the left side 1101 in fig. 11. In the example of fig. 11, it is further assumed, that a service 1104 outputs an instruction causing a response 1105. This response, in turn, triggers the processing of the rules of the rule modules RM A (111b) through RM E (1115b)

15 which are related to the processing points P0 and P4-P6, as indicated on the right side 1102 of fig. 11, starting with the service(s) 1106, and resulting in an outgoing message 1108, e.g. a generated provisional response. The processing of the services on the left side 1101 further

20 results in an outgoing message 1107, e.g. an INVITE proxy request.

Fig. 12 illustrates a hierarchical rule module processing according to a preferred embodiment of the invention.

25 When the rule engine 1200 receives an event 1201, it identifies the relevant rule modules 1216-1219 of different priorities for this event in the rule base. The rule engine 1200 generates a current event context 1202, initiates processing of the highest priority rule module

30 1216 and passes control over the generated current event context 1202 to that rule module 1216. When the processing of the highest priority rule module 1216 is completed, the rule engine 1200 invokes the subsequent rule module 1217 according to rule module priority and

35 passes the modified current event context 1206 to the next rule module 1214. Similarly, the subsequent rule

modules 1218-1219 are invoked and control over the respective current event contexts 1207 and 1214 is passed to them. Finally, the resulting event context 1215 is returned to the rule engine 1200. According to this embodiment of the invention, each invoked rule module may subsequently invoke one or more other rule modules. According to a preferred embodiment, rule modules are divided into two types of rule modules: Rule modules 1216-1219 which may be invoked by the rule engine 1200 according to their order priority, and rule modules 1220-1225 which may only be invoked by other rule modules. The latter type of rule modules may formally be identified by a special order priority, e.g. the ordering priority 0. According to this embodiment, rule modules with this priority have special restrictions associated with them: Rule modules with this ordering priority should not be invoked by the rule base processing procedure of the rule engine 1200, but are invoked from another rule module with explicit privileges to do so. Rule modules with ordering priority different from 0 are not invoked from another rule module, but only by the rule engine. Rule modules with the same owner may have the same ordering priority (different from 0), but in that case, they should have different event contexts, e.g. SIP, HTTP, etc. This means that they should not be invoked on the same event context. Rule modules with different owners may have the same ordering priority, but in that case they should not be invoked on the same property in the same event. In other words, rule modules that can be invoked on the same message property from a given event context should have different priorities, unless it is zero. The first rule module 1216 is invoked by the rule engine, before any rule modules with ordering priority 0 are invoked. This rule module is termed the root rule module. When the root rule module 1216 invokes priority 0 rule modules, they may again invoke other priority 0 rule

modules. These relationships between rule modules is termed the rule module hierarchy. This mechanism of invoking rule modules from within rule modules gives a hierarchical distribution of administrative domains, also called the hierarchical administrative domain model. The benefits with this model is that it is easy to administer, and gives a great amount of control to the master rule modules. It is a further advantage that additional rule modules may be added to an existing hierarchy without having to reassign a large number of priorities to subsequent rule modules.

The above rules avoid that rule modules are invoked twice by accident in the rule base processing procedure.

In the example illustrated in fig. 12, rule module 1216 invokes rule modules 1221 and 1222 in that order, passes control of the corresponding event contexts 1203-1204, and receives the resulting current event context 1205. Similarly, rule module 1218 invokes rule modules 1222-1223 with the corresponding event contexts 1208-1209 and the resulting event context 1213. Finally, the rule module 1223 invoked by rule module 1218 further invokes rule modules 1224-1225 with the corresponding event contexts 1210-1211 and the resulting event context 1212.

Hence, a hierarchical invocation process is performed, generating a sequence of current event contexts indicated by the enumeration of the event contexts 1202-1215.

It is noted that, if there are more than one processing points defined, the above description of rule module processing should be understood as per processing point, i.e. for each processing point a corresponding hierarchy of rule modules is processed.

An important task of the rule base processing procedure of the rule engine 1200 is the mapping of SIP events to relevant rule modules. This process very much depends on the contractual relationships defined on the domain in question. A SIP event includes a number of message properties which can be used to detect a possible contractual relationship. These properties include siprequest.from, siprequest.to, siprequest.RequestURI, sipresponse.from, sipresponse.to, sipresponse.contact, and sipresponse.via. In particular, the From and Request-URI of a SIP message are important properties of an event used to detect a contractual relationship with a network operator or 3rd party service provider. However, other message properties may be taken into account, e.g. contact headers and via headers. A message property that can be used to detect a contractual relationship and thus trigger a rule module will be referred to as a rule module trigger (RMTrigger).

When a SIP request message arrives at a SIP Server, at least one of the From or Request-URI message properties should specify a subscriber that has a contractual relationship with one of the network operators of the SIP Server. The message properties of From or the Request-URI may contain a domain name, like netopX.com, or an IP address of one of the network operators of the SIP Server, like 123.123.123.000. In addition, the From and Request-URI include a subscriber identifier, like a name or phone number. This is contained in the SIP URL or TEL URL. The SIP-URL also has parameters like "transport-param", "user-param" and "other-param". These parameters may include information specific to a network operator, as the IMSI to identify mobile terminal subscribers and terminals.

For example, using the above properties a SIP event may be mapped to a unique network operator, if a number of requirements are fulfilled, for example, if each of the network operators of the SIP Server has a unique domain name, like netopX.com and netopY.com, and if they have different IP addresses. If they share the same IP address, conjectural relationship resolutions may be ambiguous, as the message properties of the From and Request-URI may contain an IP address and not a domain name.

Third party service providers that upload/register service applications on the SIP server have a contractual relationship with the network operator which can be detected. For example, the 3rd party service provider may receive an allocated ID from the network operator with whom the 3rd party service provider has a contractual relationship. If the network operator has a domain name, e.g. netopX.com, and associated IP address, the 3rd party service provider with unique name partyZ, may receive allocated IDs following a name convention, e.g. "partyZ.netopX.com". Now the From and Request-URI may be matched to relevant rule modules. As netopX.com and partyZ.netopX.com both should be invoked on an event containing netopX.com in either the From or Request-URI, they should have different explicit ordering priorities. It is understood, that other name conventions may be introduced instead.

It is understood, that the mechanisms of order priority and rule module hierarchy may be applied independently of each other or in combination with each other, as was described in connection with fig. 12.

Fig. 13 shows an example of the flow of SIP messages and instruction sets according to an embodiment of the

invention. The message parser 1304 converts the original SIP message 1306 into an original SIP message object 1303. This is used to invoke the initial rule module 1301 and, hence, the first sequence of service applications.

5 When a rule module 1301 is loaded into the rule engine 1302 and executed, one or more actions are reached, depending on the signalling message object 1303 which was received from the message parser 1304 and on the basis of which rule module 1301 was identified, e.g. as described
10 in connection with fig. 12. When an action is reached, the service application 1305 associated with the action is invoked, that is, it is activated, possibly with parameters, and possibly with access to the entire signalling message object 1303 which triggered the
15 application. The access privileges to the entire signalling message depend on the functionality and scope of the used service access API, e.g. CGI API versus OSA API, as well as the privileges of the owner of the rule module 1301 and the privileges of the invoked service
20 application 1305. The initial signalling message object 1303 represents the entire set of message properties embedded in the original SIP message 1306 which may include multiple media types. The invoked service application 1305 performs some processing based on the
25 signalling message 1303 and hands back the result 1307 to the rule engine 1302. This result is referred to as the instruction set, as it can potentially contain multiple instructions, e.g. CGI instructions, OSA API instructions etc. When multiple services are invoked based on one
30 signalling Message, the rule module processor will end up with multiple instruction sets. This set of instruction sets is termed the "instruction set base".

The rule engine 1302 and the application execution engine
35 determine which instructions to mediate to the SIP Server default behaviour 1308. The instruction set may be

filtered based on privileges and feature interaction resolution, before mediating it to the SIP server default behaviour 1308 which may further merge the instruction set with a default SIP instruction set 1312. The
5 resulting set of instructions is termed the "resulting instruction set". The Resulting SIP message object(s) 1309 represent(s) the actual SIP messages, which are to be sent downstream or upstream. Hence, this leads to one or more resulting SIP message objects 1309 which are sent
10 to the message converter 1310 which, subsequently, transforms it into one or more SIP messages 1311 to be sent.

Fig. 14 illustrates a mechanism for managing multiple
15 instruction sets according to a preferred embodiment of the invention. According to this embodiment, an event context is a representation of a SIP message object, but it may contain additional information, such as feature interaction information as well as the MIME types
20 included in the message body of the SIP message. The original event context 1401 represents the original SIP message object 1402. The current event context 1403a-b is an event context based on the instruction set generated by the service applications 1404-1405. At any given time
25 one service application is in control of the current event context. Even when multiple service applications are concurrently running and applying control to the handling of the transaction, only one of them has the control of the current event context at any time. In the
30 example of fig. 14, two service applications 1404-1405 are shown. Initially, service application 1404 has control over the current event context 1403a. When the application 1404 has completed its processing, the control over the current event context is handed over to
35 the service application 1405. At this point, service application 1405 controls the current event context

1403b, i.e. it has the right to read and write it. In one embodiment of the invention other service applications may have the right to read the current event context 1403b.

5

Hence, triggering of subsequent service applications is based on the current event context. The current event context is further cleaned for possible feature interaction problems, so subsequently invoked service applications can rely on it. When control is handed over from one service application to the next through the invocation mechanism provided by the rule engine, the ownership of the current event context is handed over. In this way, when all service applications have been invoked and applied their instructions, the final resulting event context 1406 is achieved. Thus, the resulting event context 1406 represents the resulting SIP message object(s) 1407. Note that this mechanism does not exclude parallel processing, since the service applications 1404 and 1405 may be executed in parallel. However, according to this embodiment, a sequence of updating the current event context is enforced.

It is further noted that a service application may fork a request to multiple destinations. This results in the generation of multiple current event contexts.

Fig. 15 illustrates a tree of cascaded chains of service applications according to an embodiment of the invention. When services are executed on different SIP servers and apply control to a session, they provide a natural ordering of services, even if they are unaware of each other. This ordering may be said to be upstream or downstream. A downstream ordering is the ordering of services as they are invoked downstream from the origin client to the destination server. An upstream ordering is

the ordering of services as they are invoked upstream from the destination server to the origin client. When services which apply control to one instance of a session are invoked in a given order on the same SIP server, the
5 services can be thought of as cascaded according to the same ordering principle. This ordering principle is called cascading of services and the chain of services is called the cascaded chain of services.

10 When a request is received, the earlier a service is invoked based on this event, the more logically upstream it is considered to be in the chain of cascaded services. When a response is received then the earlier the service is invoked, based on this event, the more logically
15 downstream it is considered to be in the chain of cascaded services. Hence, the applications are treated as if they were triggered on different hosts.

This model is conceptually simple and provides a natural
20 algorithm for resolving conflicts between the instructions of multiple service applications at the same SIP event.

On reception of a SIP event the actions are executed in
25 order of priority in the following manner:

- 1) Control is passed to the first application.
- 2) Some response is received from the first application
- 3) Control is passed to the second application
- 4) Some response is received from the second
30 application and so on.

However, if the first application terminates the request, the second application is not invoked.

In this way a decision about whether to invoke a
35 subsequent application can depend on the output from the previous application. Furthermore, instruction priorities

associated with actions are by default ordered according to the cascading principle.

If an application forks a request, there is not a simple chain of cascaded services but rather a tree of cascaded services. This will be referred to as the "request tree". A request tree represents a number of cascaded chains of applications. Each path from the root of the tree to one of the leaves represents a cascaded chain.

Fig. 15 shows an example of a request tree where a service application APP1 is invoked by the rule module execution module 1501 with the original event context OEC as an input. When the control is handed over from the application APP1 to the subsequent application APP2, the application APP2 gains control over the current event context EC2. The application APP2 forks the request creating event context EC3 and event context EC4. A lower priority action invokes application APP3 due to one or more properties in event context EC3. Another lower priority action invokes another application APP4 due to one or more properties in event context EC4. This leads to a tree-like structure representing the trail of invoked applications. In this tree the branches are representations of event contexts. The nodes are representations of triggers. The root of the tree is the original event context OEC. The leaves of the tree are the resulting event contexts REC1 and REC2.

When there are no more actions in a rule module all the current event contexts are fed back to the rule base processing procedure 1502. This procedure may invoke another rule module which will construct more of the request tree. Consequently, the tree may become quite large with many nodes and branches depending on the services that are triggered.

It is noted that in some cases an application may send an asynchronous feature instruction to the service support environment, i.e. not related to an existing transaction.

5 In this case the application has started a new transaction and is considered to be at the root of a new cascaded chain tree.

In order to allow parallel processing of services, the
10 service support environment may be capable of passing control to more than one service simultaneously. This is not in contradiction with the cascaded services model, because the instructions from the parallel invoked services should be applied in the order of the cascading.
15 If the more downstream service responds before the more upstream service, the rule engine waits for the more upstream service to respond, before it can mediate the instructions. The instructions are mediated as if the services were invoked one-by-one. The administrator
20 should be capable of specifying whether a group of actions should be applied simultaneously in this manner.

Fig. 16 shows the software components of a rule engine according to an embodiment of the invention. The rule
25 base processing procedure 1601 invokes the rule modules 1602 stored in the rule base 1603 in the right order. The rule module processing procedure 1604 executes the rules in the rule module. The service interaction module 1605 covers a set of functions 1700 including enforcement of
30 the triggering, feature interaction, privileges and rights. The functions 1700 will be described in greater detail in connection with fig. 17.

When a SIP event is reported to the rule engine in the
35 form of an original event context 1609, the rule base processing procedure 1601 is executed in order to find

and execute the correct rule modules in the correct order. The rule base processing procedure 1601 passes the rule modules to be processed and the original event context 1610 to the rule module processing procedure 1604. The ordering of the rule modules together with the ordering of the rules within each rule module determines the ordering of patterns and actions 1606-1608 processed by the rule module processing procedure 1604. The actions invoke corresponding service applications 1611-1613. When invoking service application 1611, the original event context is passed as the current event context CEC1 to the service interaction module 1605 which, in turn, passes the event context EC1 to the service application 1611. The service application 1611 results in a set of feature instructions 1614 which cause an update of the current event context by the service interaction module 1605. The resulting current event context 1615 is returned to the rule module processing procedure 1604. The service applications 1612 and 1613 are invoked in a similar manner, resulting in corresponding instruction sets 1616-1617 which are converted by the service interaction module 1605 into corresponding current event contexts 1618-1619. During this conversion, the service interaction module 1605 performs authorization checks and, if a service application 1613 returns an unauthorized instruction set 1617, the unauthorized instructions are not converted. Hence, the current event context 1619 does not differ from the incoming event context 1620. Preferably, the service application is notified 1621 about this failure. It is noted that a service may result in the spawning of multiple current event contexts as described in connection with figs. 15 and 18. When a rule module is finished processing the patterns and actions 1606-1608, the final set of current event contexts 1622 is returned to the rule base processing procedure 1601. Subsequent invocations of rule

modules will be dependent on this set of current event contexts. It is noted that a service application may arm/disarm events and triggers for future events, as will be described in greater detail below. Hence, a service application may further return an arming request 1623.

Fig. 17 shows steps performed by the service interaction module between the processing of the rule module and the processing of the service application in the embodiment of fig. 16. This functionality 1700 includes feature interaction management and checking of privileges performed by the rule engine 1702 and the application execution engine 1703. When the invocation command and the current event context CECa is received from the rule module processor 1601, the privileges of the rule module are checked in step 1704. The privileges of the rule module are specified in an access control list ACL2 associated with the rule module. If the rule module has the privileges to issue invocation commands, in step 1705, the rule engine 1702 sends the invocation command to the application execution engine 1703 via the application execution engine manager (not shown). It may only be necessary to send part of the current event context to the application execution engine 1703, as denoted by $f(CECa)$.

In step 1706, the application execution engine 1703 converts the received event context $f(CECa)$ into a suitable data format for the service application 1707 to be invoked. This may include the conversion to a suitable name space for the invocation of that service application. Furthermore, in step 1708, the application execution engine 1703 checks privileges and rights. An access control list ACL3 associated with the service application 1707 may specify the privileges of the service application 1707 to access the value of $f(CECa)$

or part of it. Furthermore, the rule module has access to the service application 1707 depending on another access control list ACL4, also associated with the service application 1707. If access is granted, the service application 1707 is invoked in step 1709, and relevant parts of the event context g(ECa) are provided as an input. Subsequently, the service application 1707 returns control to the rule engine 1702 via the application execution engine 1703 by sending the instruction set 1710 or an internal representation of the instruction set. In step 1711, privileges and rights of the service application 1707 to issue these instructions are checked, e.g. based on an access control list ACL5 associated to the service application 1707. If the service application 1707 has the privileges to issue these instructions, the instructions are converted from the internal representation (step 1712), and the converted instructions 1713 are forwarded to the rule engine. Furthermore, any arming requests 1716 are forwarded as well. In step 1714, possible feature interaction problems with earlier issued instructions from previously authorized and invoked service applications are resolved. This resolution is based on whether previously invoked applications have protected message properties in the current event context. The next current event context CECb is generated and returned to the rule module processor 1601. Finally, in step 1715, the rule engine stores information about the invoked service in memory, e.g. by building the request tree described in connection with fig. 14. The request tree is used in connection with the reporting of future events.

As mentioned above, some services, e.g. monitoring applications, are interested in future events. In order to tell the rule engine about this interest in future events, the service application has to request for event

reports, also termed dynamically arming for event reports.

5 According to a preferred embodiment of the invention, dynamically arming of transaction events is tied to the processing of a SIP transaction, which is bounded in time. Transactions may typically last from between some milliseconds to a few minutes, depending on the configuration of the SIP server. Since, dynamically
10 arming of transaction events only applies to the lifetime of a transaction, this type of arming is non-permanent and implicitly disarmed when the transaction is ended, thereby providing a fast mechanism. If an application arms events which pertain to the transaction it was
15 triggered in, the application maintains its place in the cascaded services model. Responses are reported to all the applications which have armed them, starting with the leaf of the tree, i.e. the downstream most application. During the lifetime of a transaction the request tree
20 exists in the memory of the rule engine. Subsequent events related to the same transaction are thus related to the request tree. Events relating to the request tree can come from applications, from downstream servers and from upstream clients.

25 Events from upstream clients enter the request tree at the root. This means that they are first reported to the application at the root of the tree. This application may terminate or redirect the event in which case it will not
30 be sent further in the original tree but a new request tree may be constructed. A request may also need to be forwarded to all the same destinations as the original request, or the request may need to be sent to one of the destinations to which the original request was forwarded.
35 Events from applications enter the request tree at the

appropriate node. Events from downstream servers enter the request tree at one of the leaves.

5 This leads to the mechanism of request-tree traversal, in which the rule engine remembers the order in which applications were invoked, i.e. the request tree. It is an advantage of this embodiment that it provides a clear rule for the order in which events are reported within the context of a transaction. The clear rule is the
10 cascaded services model which is a conceptually simple rule that can be easily understood by administrators and application designers. This also simplifies the API via which the application execution engines add rules to the rule base. Rather than having to arm for CANCEL of a
15 specific branch in a rule module it can simply arm for cancel of the given call leg.

There are various ways to implement the request tree traversal.

20

For example, the Rule engine may achieve this traversal by using the VIA header and branch parameter. In effect this would mean creating a separate instance of the rule engine every time the destination address is changed by
25 forwarding the request to the SIP server. The DNS lookup done by the SIP stack would then determine whether or not the request comes back to be processed by the server a second time. When a request is received from the SIP server it would be treated as a new Request. The rule
30 engine should have a mechanism for ensuring that services, for example those triggered on the FROM field, would not be erroneously invoked again.

Alternatively, in a preferred embodiment, the event is
35 handled at the host for the entire request tree, i.e. including all forwardings until all the leaves of the

request tree represent destinations elsewhere in the network. This embodiment has the advantage of being more efficient as fewer instances of signalling state machines and manager classes are required. For example, the

5 request tree may be implemented, by having, for each trigger of an application, an instantiation of a node object which can have pointers to the previous and the next nodes in the tree. The following pseudo code

10 fragment shows how the building of a request tree may be incorporated in the algorithms for rule base processing and rule module processing. The idea is to associate every event context with a RequestTreeNode. This is the RequestTreeNode at which the event context was created. Each node is either the root of the tree or it is

15 associated with an action execution:

```
EC[] RuleBaseManager(EC){
    OEC=EC
    Root = new RequestTreeNode
20    OEC.SourceNode = Root
    CEC[] = do OriginatingServices(OEC)
    CEC[] = do TerminatingServices(CEC[])
    REC[] = do ForwardedByServices(CEC[])
    Return REC[]
25 }

EC[] processAction(Action, EC, ...){
    ActionNode = new RequestTreeNode
    EC.SourceNode.next = ActionNode
30    ActionNode.previous = EC.SourceNode
    if RuleModulePrivilegesOK
        EC[] = invokeService(Action , EC)
        for all EC[i] do
            EC.SourceNode = ActionNode
35    endfor
    InstructionConflictResolution(EC[])
```

```
    Return EC[]  
    else  
        Return some failure indication  
}
```

5

Considering the model of a logical cascading of service applications the following two general rules may be formulated for the service execution environment to follow:

- 10 - Events travelling upstream should be reported to the logically downstream most applications first.
- Events travelling downstream should be reported to the logically upstream most applications first.

The logical cascaded order of services is kept when
15 distributing event notifications to prevent the highly complex and unmanageable situation that events can be reported to cascaded services in any order.

Preferably, the applications that arm such events are
20 invoked at processing point 1 or 4 and their sub-processing points. Preferably, the network operator may specify in the rule list for processing point 4 a point where the dynamically armed triggers should be reported. This means that the logically cascaded chain of services
25 established for the request is kept and new services are triggered either before or after this chain. The same rule can apply for subsequent requests that are related to an existing transaction. This time it will be at processing point 1 that the administrator may specify
30 when to report dynamically armed events.

Alternatively or additionally, another type of arming may be employed which will be referred to as dynamically armed triggers which are added as rules to an appropriate
35 rule module in the rule base. Dynamically armed triggers provide a less expensive mechanism in terms of processing

power and execution memory for the reporting of events which do not pertain to the transaction in which the application was triggered. Preferably, these requests for event reports should be stored in persistent storage, i.e. they become rules in the rule base. These rules can be non-permanent by specifying an expiration timer, or permanent. In the latter case the service application should explicitly disarm the request for event reports to remove it. Alternatively, they may also be armed in a 'report-once then disarm' mode. Furthermore, if no expiration time is given, a default time, e.g. 1 hour, may be applied to the rule. When this time has expired, the rule is deleted. This has the advantage of avoiding that that the server runs out of data storage capacity.

Preferably, trigger rules are added to a rule module which the application has the privileges and rights to update. It will normally be the same rule module from which the application was triggered. Exactly where in the priority order within a rule module these rules should be added may be determined by an implicit rule priority order, i.e. an integer representing where the rule should be placed within the existing rule module. When a SERL script is first added, the rules are ordered in the order they appear in the script. If there are N rules, the integers 1 through N are implicitly associated with the rules. A rule can be deleted by referring to these numbers. A rule can be added by referring to the rule that the new rule should be placed after. This has the advantage, that the amount of data that needs to be exchanged when a rule is added to a large rule module is reduced.

If no position is specified, then the rule module engine may follow the following algorithm when adding these triggers: Search the rule module for the same property

pattern as in the new rule. If a pattern is found, search for any sub-pattern, and so on. If no sub-pattern is found, insert the rule as the highest priority rule in the enclosed list of actions. If no similar pattern is in the rule module already, insert as the first pattern. This algorithm provides a logically advantageous placement of the rule. For example, it ensures that the actions for the caller where TARGET= FROM are not mixed with actions for callees where TARGET=RequestURI. It also means that rules added last are also the first encountered when an event happens. This is the simplest default behaviour.

It may be possible to indicate whether a trigger rule is permanent or should be automatically deleted once the event is reported. Which type of rules an application can add dynamically may be linked to the privileges and rights assigned to the application.

In another alternative embodiment the cascaded request tree is only maintained for the lifetime of an event. In this case, when the last application in the chain has finished its processing and the SIP event is sent then the cascaded chain is no longer relevant. Until that point it may be useful for the Rule engine to hold a representation of the cascaded chain in memory. This is because applications may run on separate servers and may take some time to respond to the invocation. While the rule Engine is waiting for a response an earlier application in the chain may cancel the request. If this happens, the rule engine will need to inform the application execution engine of the application it is waiting for. This means that the rule engine should remember which was the next application in the chain. By the time the SIP event has been forwarded all the applications in the chain should have armed rules for the

events which interest them, thus there is no longer a need for the rule engine to remember the cascaded chain. The advantage of this solution is that the rule engine itself is simple. It just reports events based on Rule
5 modules at the relevant processing points. However, this solution is complex to administrate and places the complexity of ordering of the events into the applications. i.e. the applications need to decide in which priority position the rule should be added in a
10 rule module.

Fig. 18 illustrates the tree structure of the processing of rule modules according to an embodiment of the invention. When an event is received, the rule base
15 processing unfolds a processing pattern which will be referred to as the rule base tree. When a request arrives at the SIP Server corresponding to an original event context 1801, the associated subscribers may be the originating party (i.e. caller) and/or the terminating
20 party (i.e. callee). The caller is identified by the From header field. The Callee (or current callee) is identified by the Request-URI. The From header field and the Request-URI should uniquely identify a subscriber at the SIP Server where these subscribers have a contractual
25 relationship with the network operator and service providers. Following the cascading principle, the originating services 1802 are invoked before terminating services 1803, even if the originating and terminating parties are located at the same host. A third category of
30 services may also be invoked. These are termed the forwarded-by services (not shown). This category of services is invoked, if a request is forwarded to a new destination belonging to another subscriber. In this case originating services may need to be invoked on behalf of
35 the callee. For each party, the services are invoked based on location of the rule modules as they are placed

in the different processing points 1804-1809. For each processing point, the sets of rule module priorities 1810-1814 are examined for a match. Within each priority at most one rule module 1815-1817, respectively, is
5 invoked. However, a rule module 1817 may be the root of a rule module hierarchy 1818 as described in connection with fig. 12. For simplicity, such a rule module hierarchy may be considered as a single rule module. Hence, the function of the rule module processing 1601,
10 the functions 1700 of the service interaction module and the service applications 1819 invoked by the rule modules, may be considered as a leave in the rule base tree. Based on the current event context CEC, the rule module 1817 is invoked and it and returns a resulting
15 event context REC. This resulting event context REC is considered the current event context CEC for the next rule module that is invoked. When all rule modules have been invoked and the last of them returned the resulting event context, then the resulting event context is the
20 set of SIP signaling messages that will be sent upstream and/or downstream, as answer to the original incoming SIP message.

The above processing structure which is graphically
25 illustrated in fig. 18 may further be illustrated by the following example of a pseudo code fragment:

```
EC[] RuleBaseManager(EC) {  
    OEC=EC  
30    CEC[] = do OriginatingServices(OEC)  
    CEC[] = do TerminatingServices(CEC[])  
    REC[] = do ForwardedByServices(CEC[])  
    Return REC[]  
}  
35  
EC[] OriginatingServices(EC) {
```

```
    if Subscriber(EC.From) != UNKNOWN
        return do RuleBaseProcessing(EC, TARGET="FROM")
    else
        return EC
5    endif
}

EC[] TerminatingServices(EC[]){
    for all EC[h]
10        if Subscriber(EC[h].Request-URI) != UNKNOWN
            TEMP[] =
                do RuleBaseProcessing(EC[h], TARGET="REQUEST-URI")
            EC[h] = TEMP[]
        endif
15    endfor
    return EC[]
}
```

Where:

```
20 EC[] = EC{EC[1], ..., EC[h], ..., EC[C_MAX]}
```

and EC[h] = TEMP[] means:

```
EC[] = EC{CEC[1], ..., EC[h-1], TEMP[1], ...,
        TEMP[T_MAX], EC[h+1], ..., EC[C_MAX]}
```

25

```
EC[] ForwardedByServices(EC[]){
    for all EC[h]
        if Subscriber(EC.ForwardedBy) !=
            (UNKNOWN | NOT PRESENT)
30        TEMP[] =
            RuleBaseProcessing(EC, TARGET="FORWARDEDBY")
        EC[h] = TEMP[]
    endif
    endfor
35    return EC[]
}
```

```
EC[] RuleBaseProcessing(EC, TARGET) {  
    CEC[1] = EC  
    for all (relevant) processing points CPP[j]  
5      for all priorities l  
        RuleModule = FindRuleModule(j, l)  
        for all CEC[t]  
            TEMP[] = processRuleModule(CEC[t],  
                RuleModule, TARGET, CPP[j])  
10         CEC[t] = TEMP[]  
        endfor  
    endfor  
    endfor  
    endfor  
15    Return CEC[]  
}.
```

It is noted that a service application may change the From header field or the Request-URI. Furthermore, the resulting From header field or the resulting Request-URI may belong to another subscriber which may even be unknown to the SIP Server which processes the event. If the resulting From and/or the Request-URI is a new but known subscriber at the SIP Server, the services associated with this subscriber are invoked as well. In this case the rule base processing procedure may be invoked recursively resulting in a hierarchic structure of rule base trees.

Processing of a rule module comprises the step of taking each action in turn in priority order, evaluating whether the enclosing patterns match the current event context and if so applying the action. In the following, a method according to an embodiment of the invention is described in greater detail.

The following pseudo code fragment provides a high level description of an algorithm for processing a rule module:

```
EC[] processRuleModule(EC, RM, TARGET, CPP){
5   CEC[1] = EC
   // CPP represents the current processing point CPP
   // TARGET represents the rule module trigger
   for all ACTION[i] in RM
     for all CEC[t]
10      if CEC[t].ForwardedBy != null
         if (enclosingPatternsTrue(CEC[t], Action[I]))
            TEMP[] = processAction(ACTION[i], CEC[t],...)
            for all TEMP[h]
               if (TARGET is a REQUEST_URI AND
15                  Subscriber(TEMP[h].REQUEST-URI) !=
                     Subscriber(TARGET)
                     // i.e. for terminating services
                     and if the call has been
                     forwarded to a new user then
20                  TEMP[h].ForwardedBy = TARGET
               endif
            endfor
            CEC[t] = TEMP[]
            t= t+TEMP_MAX
25          endif
        endfor
      endfor
    endfor
    Return CEC[]
30  }.
```

Here, The enclosingPatternsTrue method returns a boolean indicating if the patterns which enclose the actions are true. The processAction method may be implemented as
35 indicated in the following pseudo code segment:

```
EC[] processAction(Action, EC, ...){  
    if RuleModulePrivilegesOK  
        EC[] = invokeService(Action , EC)  
        InstructionConflictResolution(EC[])  
5        Return EC[]  
    else  
        Return a failure indication  
},
```

10 where EC[] is the set of forwardings of the event context. If it is empty, the service has terminated the request or response.

The service application may further issue instructions
15 that are not instructions to forward the event context. These are not shown in the above algorithm. Such instructions may be processed by the rule base manager.

When executing actions in priority order, the result of
20 the action may be to change the current event context. After such an action, the processing of the current rule module continues. After executing the current rule module, further rule modules may be executed according to the rule base processing procedure. The lower priority
25 order actions are only executed, if the pattern(s) that enclose them match the new message properties, as specified by the current event context. This is illustrated by the following example of a fragment of pseudo script which describes two actions enclosed by two
30 patterns:

```
<property name=Request matches=INVITE>  
    <property name=RequestURI matches="^xcorp.com">  
        do Action CPL?user=RequestURI  
35        do Action proxybehaviour  
    </property>
```

</property>.

The two actions are enclosed by two patterns which indicate that the actions should only be applied if the request is an INVITE and the RequestURI contains the domain name xcorp.com. The first action invokes a user supplied CPL script. If the user CPL script does not change the destination of the request, the standard proxy behaviour is invoked to locate the user and proxy the request. If the CPL script changes the destination so that the destination will resolve to another host, the standard proxy behaviour need not be invoked.

If the message property which triggered the rule module changes so that it represents a new user, the processing of that rule module is stopped. The purpose of this is that rule modules can be thought of as being processed on behalf of a single user. This simplifies the rule module processing and script authoring.

Fig. 19 illustrates the recursive processing of rule modules in a situation where service applications generate new event contexts according to an embodiment of the invention. The new event contexts may change the original rule module trigger and, in this case, new rule modules may be invoked recursively. For example, a subscriber's callee preferences, e.g. a CPL script, may forward the request to a new destination associated with another subscriber of the same SIP server. In this case the forwarded-to subscriber's callee preferences, e.g. another CPL script, should be invoked as well.

As illustrated by the example of fig. 19, an initial event context 1901 generates a recursion of rule module invocations, when a SIP request message is forwarded to multiple new destinations associated with other

subscriber accounts. In the example of fig. 19, the application 1914 is triggered by the rule module 1903 on behalf of subscriber A. The rule module 1903 is triggered by the original event context 1901. The application 1914
5 generates three event contexts 1904-1906, where the rule module trigger has changed in the two event contexts 1905-1906 which are associated with new subscriber accounts. The event context 1904, is an updated event context but associated with the same subscriber whose
10 rule module initially invoked application 1914. In this case the subscriber rule modules 1907-1908 associated with the new event contexts 1905-1906, respectively, are also invoked. Rule module 1907, in turn, invokes application 1915 which generates two new event contexts
15 1909 and 1910, and so forth. Consequently, a tree processing structure is created in which the leaves correspond to the set of resulting event contexts 1910-1913. Subsequently, the resulting event contexts 1910-1913 cause the SIP server to send SIP messages upstream
20 or downstream or both. It is noted that the example of fig. 19 shows a simple case assuming that the rule modules only contain one action each.

If the five applications 1914-1918 in fig. 19 have
25 requested to be notified of subsequent replies to the resulting SIP messages, according to the cascading principle, responses are notified to the logically most downstream application first. In fig. 19 responses to the resulting event context 1910 are notified first to
30 application 1918, then to application 1915 and finally to application 1914, as indicated by the line 1919. For the resulting event context 1911 the corresponding sequence is application 1915 and application 1914, as indicated by line 1920. For the resulting event context 1912 the
35 sequence is application 1916 and then application 1914, as indicated by line 1921. Finally, for the resulting

event context 1913 the sequence is application 1917 and then application 1914.

Fig. 20 illustrates a mechanism of enforcing access control in connection with rule modules according to an embodiment of the invention. Two types of access control are enforced. Access to rule modules should only be granted to authenticated and authorized parties, and access to service features should only be granted to authenticated and authorized rule modules. According to one embodiment this access control may be implemented by using so-called access control lists (ACL). Such a list may be an XML document possibly embedded in the rule module or located in the same directory as, or otherwise associated with, the rule module. Access control lists include access control rules. The access control list may enumerate each of the individuals or groups that are granted access to the rule module. The mechanism may also be used to explicitly specify privileges and rights of the rule modules, e.g. which service features the rule module may use. Therefore, it allows the administrator to manage privileges and rights of rule modules.

Furthermore, the mechanism may be used to manage privileges and rights of service applications. If a subscriber uploads a CPL script, there may be an associated rule module that invokes that service at the right time. In this case the rule module should have explicit privileges and rights to do so. This may be managed by associating access control lists with the CPL Script. The mechanism may further be used to explicitly specify privileges and rights of the service application, e.g. which service features, APIs, etc., the service application may use. Hence, it allows the administrator to manage privileges and rights of service applications.

Referring to the enumerated circles in fig. 20, the following privileges may be checked:

- 1) An original event context 2001 is sent to the rule engine after proper authentication of subscribers (A is the caller, B the callee). Rule modules and service applications are all authenticated, and have been given privileges and rights.
- 2) The rule base processor 2002 tries to access rule modules in the rule base 2003, which are associated with authenticated subscriber A or B or both. The rule base processor 2002 should not be allowed to access rule modules associated with other subscribers when processing this event. The rule base 2003 may be located on a remote server, in which case the rule engine should authenticate itself.
- 3) If allowed according to the access control list 2005, the rule base processor 2002 can invoke the loaded rule module 2004, say rule module 1 owned by A.
- 4) If allowed according to the access control list 2006, rule module 2004 may attempt to invoke another rule module 2007 associated with another owner B.
- 5) If allowed according to the access control list 2008, the rule base processor 2002 may invoke the loaded rule module 2007.
- 6) If allowed according to the access control lists 2005 and 2012, rule module 2004 may attempt to invoke service application 2010 if allowed. The service application 2010 may access the event context 2013 according to access control list 2011.
- 7) The service application 2010 may return a set of instructions 2014 which is mediated back to the rule module processor, if allowed according to access control list 2015.

8) The service application 2010 may arm/disarm for events, if allowed according to access control list 2015.

5 9) The service application 2010 may attempt to give instructions to another rule module 2007 than the one that invoked it, if allowed according to access control List 2012.

10 Access control to rule modules may be enforces by the policy nodes in a rule module. An example of a rule module access control list embedded in a rule module may look like this:

```
15 <policy>
    <rmacl>
        <rule id="rule1">
            <acl>
                <subject><subscriber>Alice</subscriber></subject>
                <privilege>
20         <read>no</read>
                <write>no</write>
                <invoke>yes</invoke>
                </privilege>
            </acl>
25     </rule>
    </rmacl>
</policy>
```

30 This policy could be applied to the entire rule module, in which case it states that the subscriber Alice, may invoke this rule module, but she must not read or write it.

35 Such policy XML scripts may be embedded within rule modules, but they may also be associated data elements. For example, the network operator may link a policy XML

script to the rule module of a subscriber, specifying the privileges granted by the network operator to the owner of the rule module. These privileges could specify the services or service features the rule module is allowed to invoke, as illustrated by the following example of such a service access control list:

```
<policy>
  <sac1>
10    <rule id="rule1">
      <acl>
        <subject><subscriber>Alice</subscriber></subject>
        <privilege>
          //this rule module may use the logging
15          function
            <service>
              <name>logging</name>
              <invoke>yes</invoke>
            </service>
20          //this rule module may invoke other rule
            modules
            <service>
              <name>RuleEngine</name>
              <invoke>yes</invoke>
25          </service>
        </privilege>
      </acl>
    </rule>
  </sac1>
30 </policy>.
```

It is noted that a SIP server service support environment may further be adapted to generate accounting records, e.g. for content charging, applications charging, usage charging, or the like. For example, the records may be provided via logs from SERL scripts or CPL scripts,

managed via library functions or a service applications, etc.

Furthermore, a system implementing service triggering
5 according to the invention, preferably implements security measures in order to preserve safety and integrity of the SIP Node even though subscribers and 3rd party service providers may be allowed to upload not only service applications, but also instructions on how and
10 when to invoke them, including some degree of feature interaction resolution. possible security measures include the configuration of privileges and rights of rule modules, service applications, name space convention policies, authentication mechanisms, authorization
15 mechanisms, access protection, authentication and validation of uploaded rule modules, logging and monitoring, etc.

It is noted that the invention has primarily been
20 described in connection with network services. However it is also applicable to be used in end user equipment.

Furthermore, it is noted that the invention, although primarily described in connection with SIP, may embrace
25 other signalling protocols as well. SERL is not limited to invoke services based on SIP events, but may invoke any type of service application based on any type of event, in the context of any type of business model. The invention may be applied to manage services for any SIP
30 enabled node. Using SERL scripts, services can be invoked from nodes that implements user agents, registrars, redirect servers or proxy servers.

Finally, it is noted that in the 3GPP architecture all
35 subscriber data is administered in the so-called Home Subscriber server (HSS). SIP related applications are

invoked from a node called a Serving Call State Control Function (S-CSCF). When a subscriber connects to a network his User Equipment (UE) performs a CSCF discovery to select an appropriate S-CSCF. The S-CSCF registers
5 with the HSS that it is serving the subscriber in question. Service triggers could then be transported from the HSS to the S-CSCF in the form of service execution rules in the SERL format. The HSS could also use the Service Execution Rules associated with a subscriber to
10 decide to allocate a different S-CSCF based on which S-CSCF has the correct services installed. Hence, an embodiment of the invention may be used in order for the HSS to place the subscriber based triggers at the correct processing point and priority and thus in the correct
15 priority order with permanent services installed at the S-CSCF. Hence, a mechanism according to the invention may be embedded in the 3GPP IPMM domain.

CLAIMS

1. A method of managing a plurality of services
(309,310,311;1305;1611,1612,1613) triggered by a message
5 (1306;1609) of a session protocol controlling a
communications session,

the method comprising the steps of

10 obtaining a number of execution rules
(308;401,402;1606,1607,1608) each of which specifying a
condition (403,404,405) for invoking a service;

processing the execution rules in a predetermined order,
15 a first execution rule (1606) causing a first service
(1611) to be invoked, if the message fulfils a first
condition, resulting in a first modified message (1615);
and a second execution rule (1607) causing a second
service (1612) to be invoked with the first modified
20 message as an input, if the first modified message
fulfils a second condition.

2. A method according to claim 1, characterised in that
said number of execution rules is grouped into a number
25 of rule modules (308;400;1216,1217,1218,1219;1301;1602),
each rule module including a number of execution rules;
and the method further comprises the steps of

- processing a first one (1216) of said number of rule
modules resulting in a first accumulatively modified
30 message (1206); and
- invoking processing a second one (1217) of said number
of rule modules providing the first accumulatively
modified message as an input.

3. A method according to claim 2, characterised in that each rule module has associated with it a priority indicative of an order of processing of said number of rule modules.

5

4. A method according to claim 2 or 3, characterised in that each rule module corresponds to a rule module owner (409) authorised to edit the rule module.

10 5. A method according to any one of claims 2 through 4, characterised in that the first rule module has assigned to it a privilege indicative of an authority to alter a lock flag related to a predetermined part of the accumulatively modified message and specifying whether
15 said predetermined part of the accumulatively modified message may be modified by services invoked from at least the second rule module.

20 6. A method according to claim 5, characterised in that the step of invoking processing the second rule module further comprises the step of setting said lock flag to prevent modification of the predetermined part of the accumulatively modified message by services invoked from the second rule module, unless the lock flag was marked
25 unset by the first rule module.

7. A method according to any one of claims 2 through 6, characterised in that the step of obtaining a number of execution rules further comprises the step of detecting a
30 predetermined contractual relationship based on header information of the message; and selecting a number of rule modules based on said detected contractual relationship.

35 8. A method according to any one of claims 2 through 7, characterised in that the step of processing the first

rule module (1216) further comprises the step of invoking a predetermined third rule module (1220).

9. A method according to any one of the claims 2 through 5 8, characterised in that the first and second rule modules are related to respective first and second access control lists (411) specifying access rights to the corresponding first or second rule module.
- 10 10. A method according to any one of the claims 2 through 9, characterised in that the first and second rule modules comprise respective first and second scripts in a predetermined mark-up language.
- 15 11. A method according to any one of claims 1 through 10, characterised in that the message comprises a first and a second set of attributes; the execution rules are grouped into at least a first and a second processing class (509,510;903,904) of execution rules according to 20 corresponding constraints, where the second processing class is restricted to only modify attributes of the second set of attributes; and the step of processing the execution rules further comprises the step of processing the execution rules of the first processing class before 25 processing any execution rule of the second processing class.
12. A method according to any one of claims 2 through 10, characterised in that the message comprises a first and a 30 second set of attributes; the execution rules are grouped into at least a first and a second processing class (903,904) of execution rules according to corresponding constraints, where the second processing class is restricted to only modify attributes of the second set of 35 attributes; and the method further comprises the step of repeating the steps of processing the first rule module

(901) and invoking processing the second rule module (902), where in each repetition the processing of the first and second rule modules is limited to execution rules of a corresponding processing class, and where each
5 repetition results in a corresponding accumulatively modified message which is used as an input for a subsequent repetition.

13. A method according to claim 11 or 12, characterised
10 in that the processing classes are defined separately for execution rules triggered by requests (607) and responses (608) of the session protocol.

14. A method according to any one of the claims 11
15 through 13, characterised in that the first set of attributes comprises signalling properties of the message.

15. A method according to any one of the claims 11
20 through 14, characterised in that the processing classes correspond to predetermined locations in a round trip message flow according to the session protocol.

16. A method according to any one of the claims 11
25 through 15, characterised in that the processing classes include a first processing class (P1) of execution rules which impact signalling properties (702) of the message, a second processing class of execution rules (P2) which impact non-signalling message body content (704) of the
30 message, and a third processing class (P3) of execution rules which neither impact the signalling properties nor the non-signalling message body content of the message.

17. A method according to claim 16, characterised in that
35 a resulting modified message (703) is generated when all

execution rules of the first and second processing classes are processed.

18. A method according to any one of the claims 1 through
5 17, characterised in that invoking the first service
further results in a second modified message; and the
method further comprises the steps of processing
subsequent execution rules with the first modified
message as an input; and processing subsequent execution
10 rules with the second modified message as an input.

19. A method according to any one of the claims 1 through
18, characterised in that the method further comprises
the steps of

15 storing information about which services are executed and
information about which order the services are executed
in;

20 receiving from the first service a request for returning
a notification to the first service, if a predetermined
event occurs;

storing the request in relation to the stored
25 information; and

upon occurrence of the event, notifying the first service
according to the stored information.

30 20. A method according to any one of the claims 2 through
19, characterised in that the execution modules comprise
computer-readable scripts, and the predetermined order of
processing the execution rules is determined by the order
of execution rules in said scripts.

35

21. A method according to any one of the claims 1 through 20, characterised in that the step of processing the execution rules further comprises the step of marking at least a part of the first modified message as being
5 locked for modification by at least the second service.

22. A method according to any one of the claims 1 through 21, characterised in that the session is related to a number of subscribers including a caller and a callee; a
10 service is adapted to be triggered by a request from a subscriber; and the method further comprises the step of invoking services requested by the caller before invoking any service requested by the callee.

23. A method according to any one of the claims 1 through 22, characterised in that the session protocol is a session initiation protocol.
15

24. A method according to any one of the claims 1 through 23, characterised in that the communications session is a multimedia session.
20

25. A method according to any one of the claims 1 through 24, characterised in that the execution rules are adapted
25 to be dynamically updated.

26. A data processing system comprising a service execution environment module (201;301) adapted to invoke a plurality of services (309,310,311) triggered by a
30 message of a session protocol controlling a communications session;

characterised in that

35 the data processing system further comprises a storage medium (307) adapted to store a plurality of execution

rules each of which specifying a condition for invoking a service; and

the service execution environment module comprises a rule engine module (303) adapted to

- retrieve a number of execution rules (308); and
- process the execution rules in a predetermined order, a first execution rule causing a first service to be invoked, if the message fulfils a first condition, resulting in a first modified message; and a second execution rule causing a second service to be invoked with the first modified message as an input, if the first modified message fulfils a second condition.

27. In a data processing system, a service execution environment module (201;301) adapted to invoke a plurality of services triggered by a message of a session protocol controlling a communications session;

characterised in that

the service execution environment module comprises a rule engine module (303) adapted to

- retrieve a number of execution rules (308) each of which specifying a condition for invoking a service; and
- process the execution rules in a predetermined order, a first execution rule causing a first service to be invoked, if the message fulfils a first condition, resulting in a first modified message; and a second execution rule causing a second service to be invoked with the first modified message as an input, if the first modified message fulfils a second condition.

28. A service execution environment module according to claim 27, characterised in that the rule engine module is

adapted to interpret a predetermined rule execution specification language.

29. A service execution environment module according to
5 claim 27 or 28, characterised in that the rule engine
module further comprises a rule base processor module
(1601) adapted to retrieve the execution rules; and a
rule module processing module (314;1604,1605) adapted to
process the execution rules.

10

30. A software program comprising code means adapted to
perform; when executed on a data processing system, the
steps of the method of any one of the claims 1 through
25.

15

31. A software program according to claim 30,
characterised in that the software program is embodied on
a computer-readable medium.

20

32. A data record comprising a rule module for use in a
method of any one of the claims 2 through 25.

25

33. A method of managing the deployment of a service in a
service network, the method comprising the steps of

specifying in a computer-readable script a number of
privileges and rights to be granted to the service during
operation;

30

analysing potential causes of feature interaction;

specifying a deployment strategy for said service as a
set of execution rules stored as a computer readable rule
module script.

35

34. A method according to claim 33, characterised in that the method further comprises the step of assigning a rule module priority to the rule module script.




Fig. 1




Fig. 2




Fig. 3

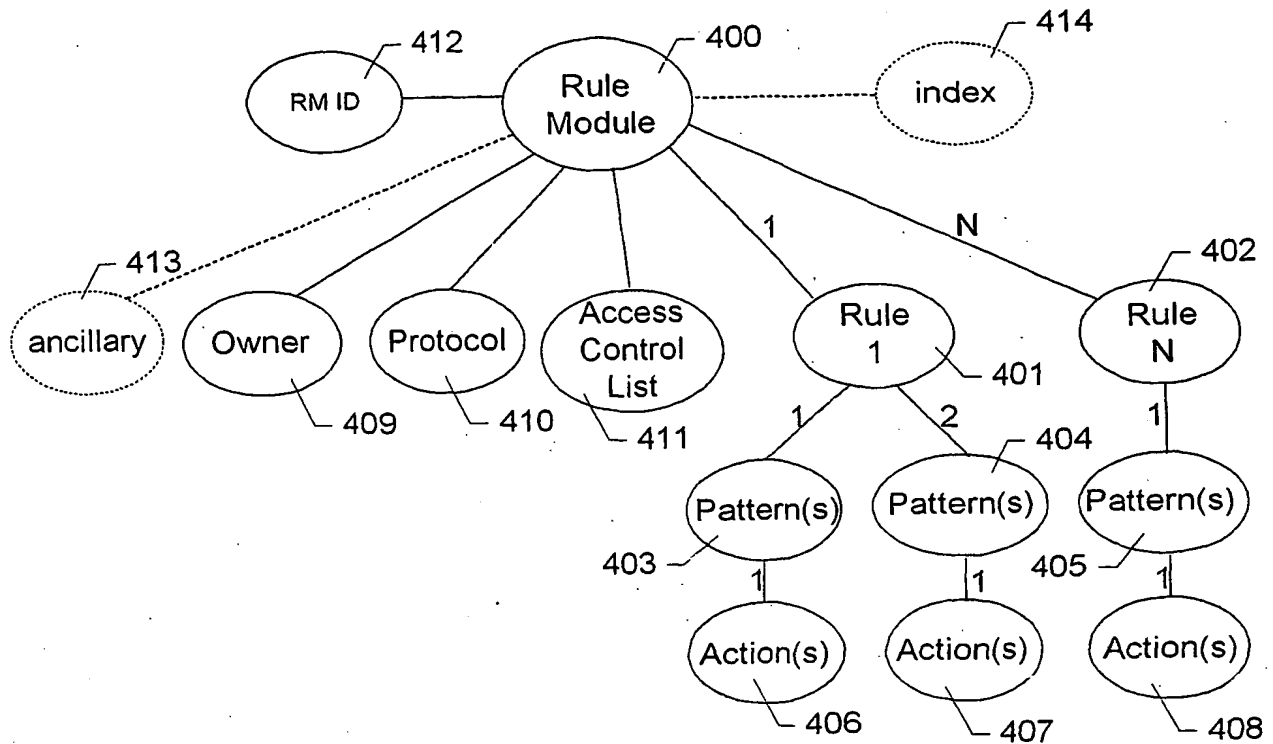


Fig. 4


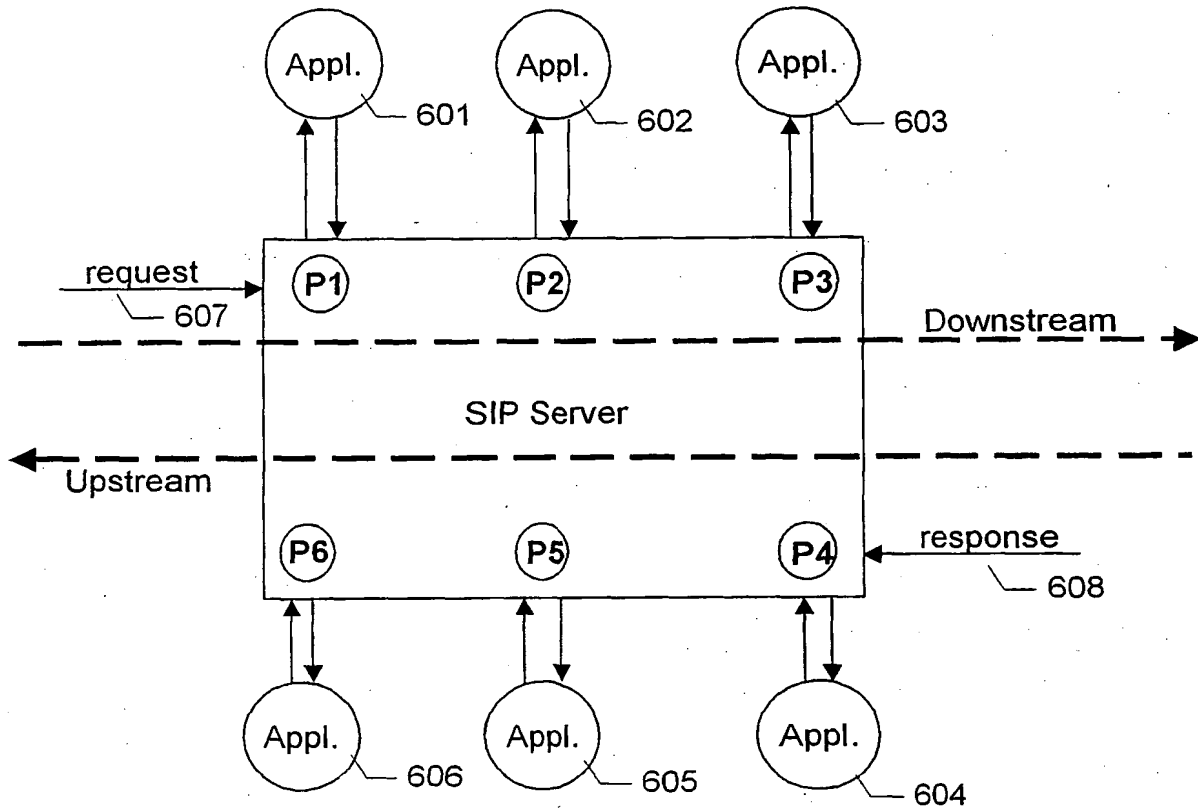



Fig. 5

**Fig. 6****Fig. 7**




Fig. 8




Fig. 9




Fig. 10




Fig. 13




Fig. 11

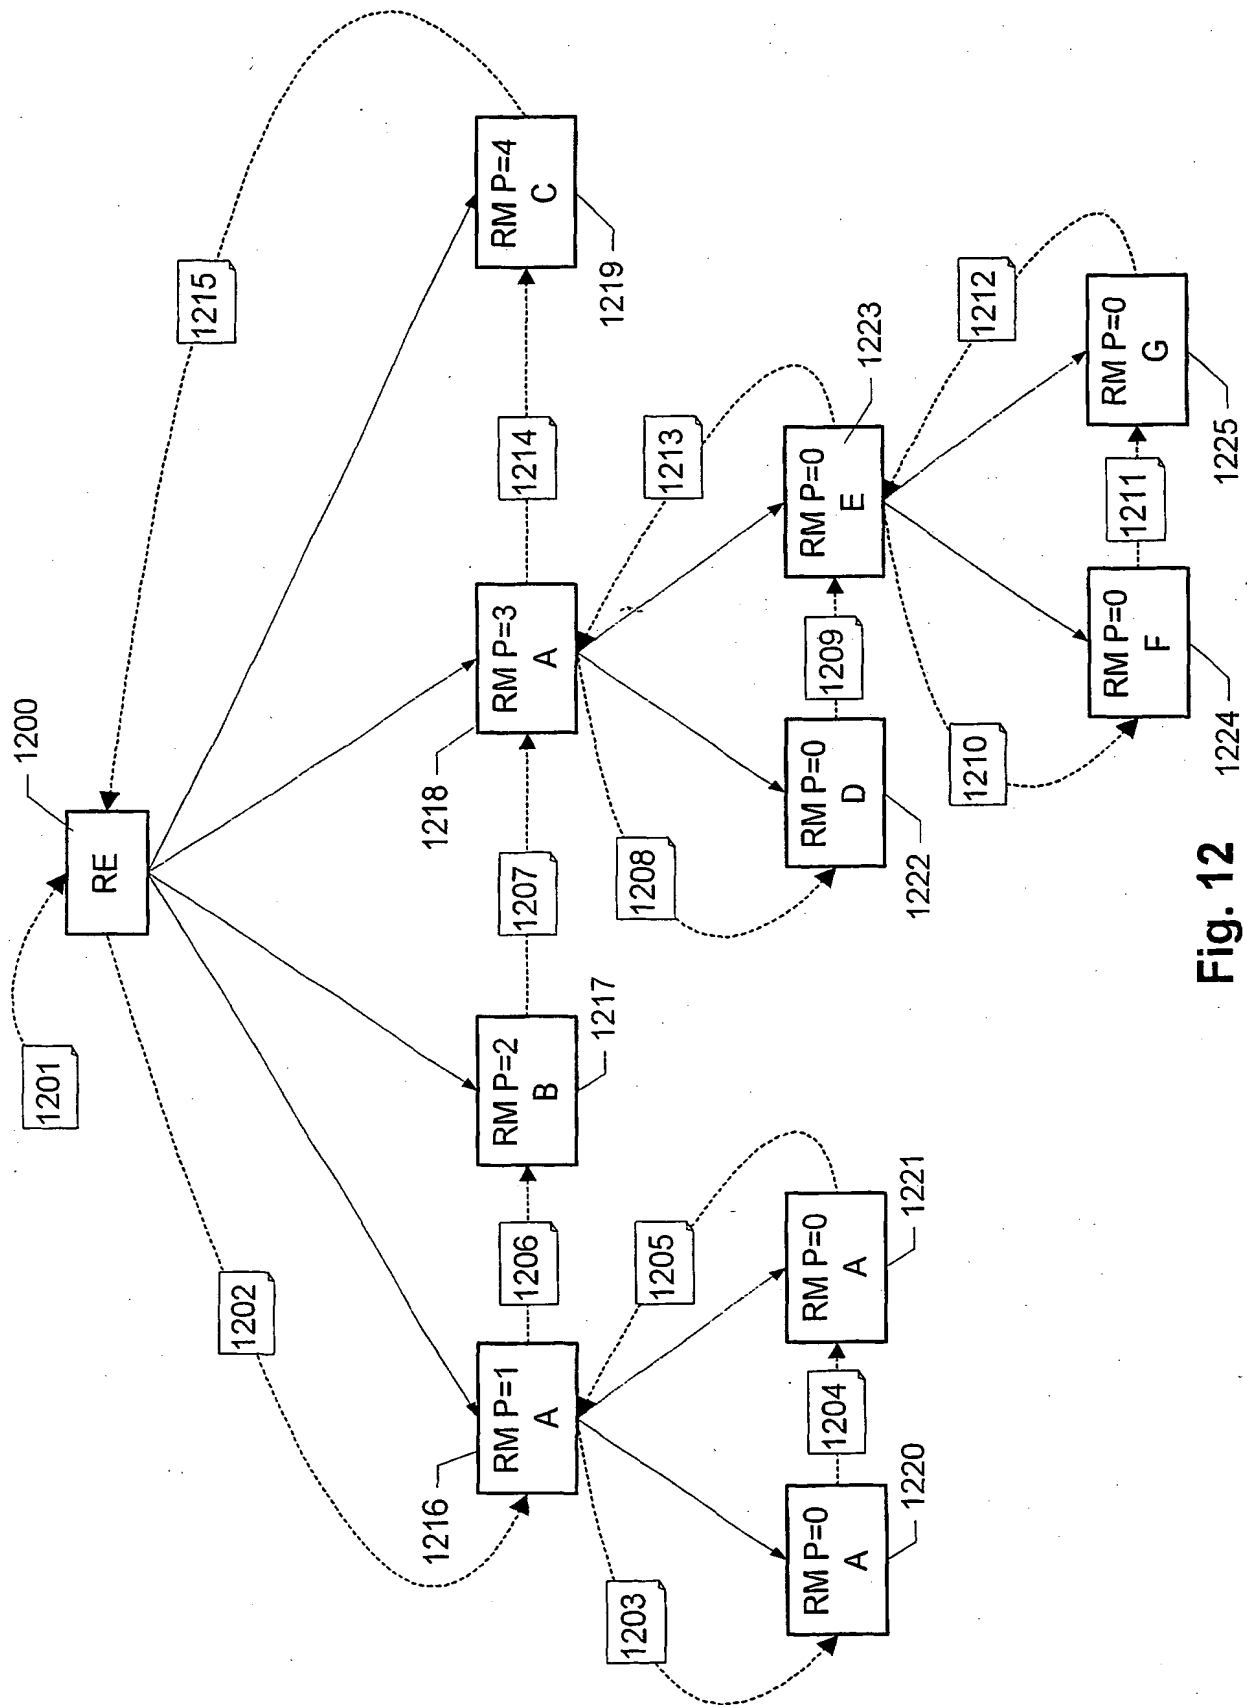


Fig. 12

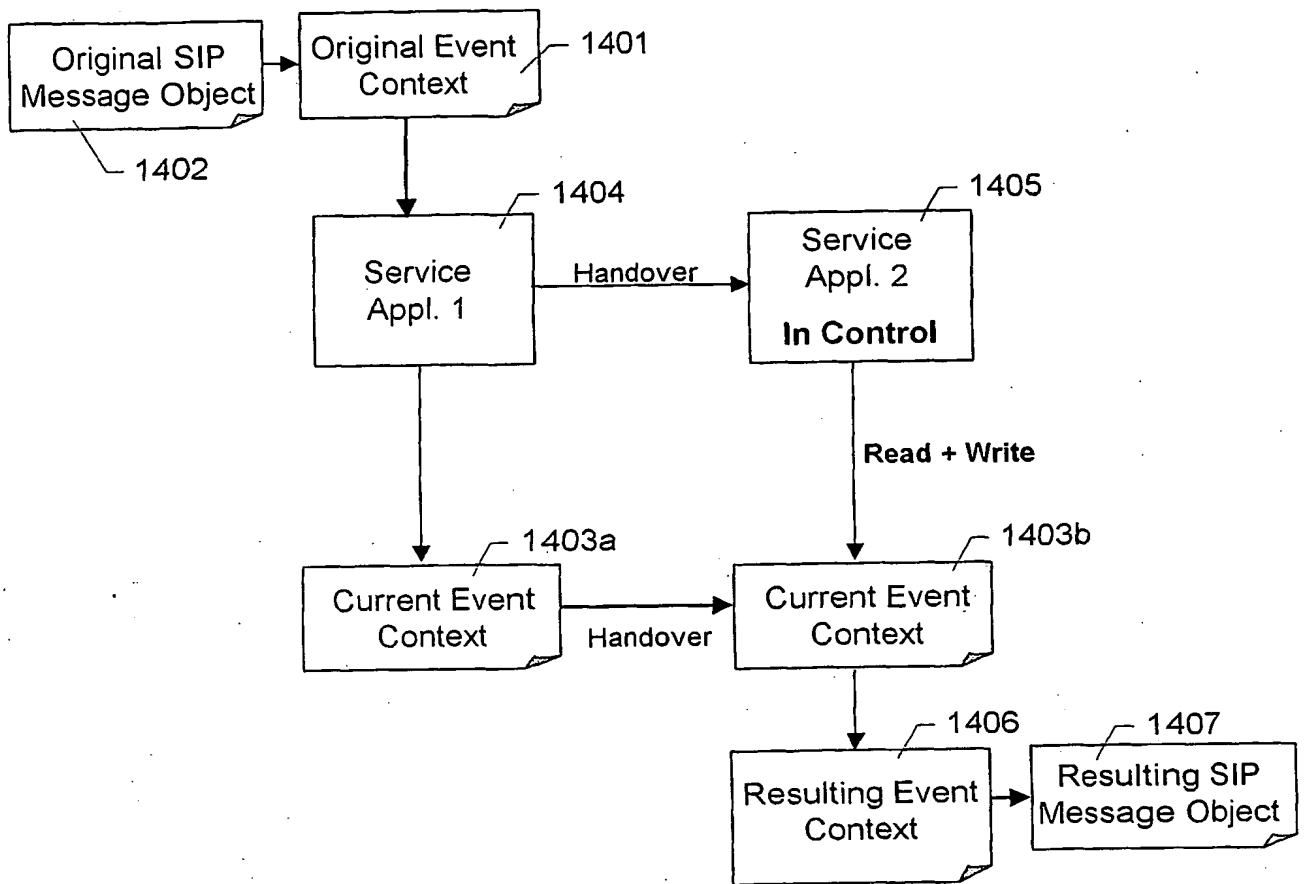


Fig. 14




Fig. 15




Fig. 16




Fig. 17




Fig. 18




Fig. 19

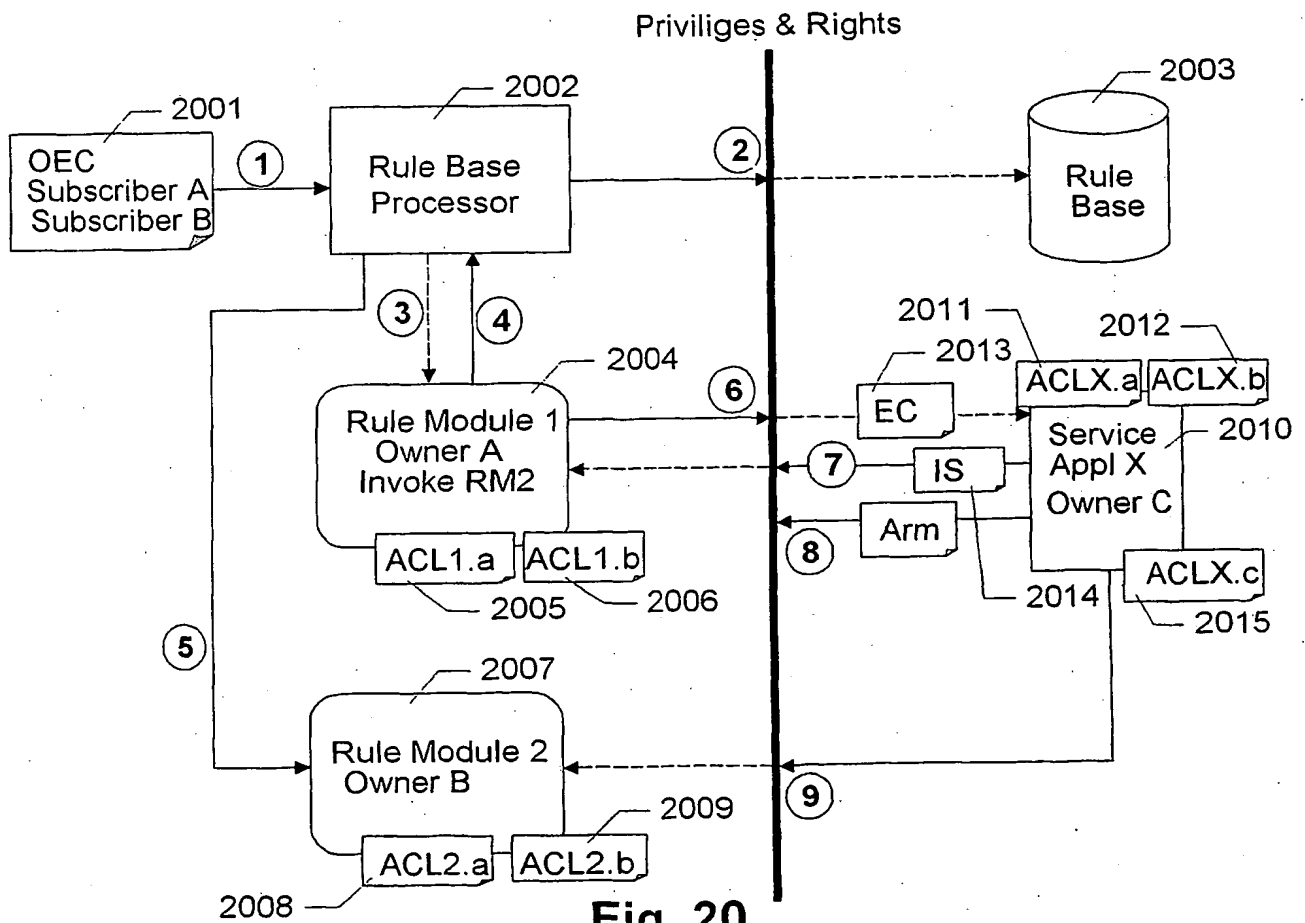


Fig. 20

INTERNATIONAL SEARCH REPORT

International Application No

PCT/EP 02/04867

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04Q3/00 H04L29/06 H04L12/24

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04Q H04M H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	ZAVE, PAMELA: "An Architecture for Three Challenging Features" INTERNET TELEPHONY WORKSHOP 2001, 'Online! 3 April 2001 (2001-04-03), XP002196887 New York Retrieved from the Internet: <URL:http://iptel.org/2001/pg/final_program/28.pdf> 'retrieved on 2002-04-22! the whole document	1-4, 7, 9, 10, 18, 22-24, 26, 27, 30, 32, 33
Y	---	19, 33
	--- -/--	

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

G document member of the same patent family

Date of the actual completion of the international search

2 July 2002

Date of mailing of the international search report

11/07/2002

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

Mannekens, J

INTERNATIONAL SEARCH REPORT

International Application No

PCT/EP 02/04867

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	KRISTENSEN, ANDERS: "Programming SIP Services" IP TELEPHONY WORKSHOP 2000, 'Online! 12 April 2000 (2000-04-12), XP002196888 Berlin Retrieved from the Internet: <URL:http://www.fokus.gmd.de/research/cc/g lone/projects/iptel2000/crc/68/68-kristens en-paper.pdf> 'retrieved on 2002-04-22!	1-4,7,9, 10,18, 22-24, 26,27, 30,32,33
A	the whole document	1-34
Y	ROSENBERG, JONATHAN: "Programming Internet Telephony Services" IEEE NETWORK MAGAZINE, 'Online! vol. 13, no. 3, 1 June 1999 (1999-06-01), pages 0-25, XP002196889 Retrieved from the Internet: <URL:http://www.jdrosen.net/papers/program _report.ps> 'retrieved on 2002-04-22!	19
A	page 6, line 43 -page 7, line 6 the whole document	1-34
Y	IBM: "Access Control List - Simple Example 1" XML SECURITY SUITE, 'Online! 14 October 2000 (2000-10-14), XP002196890 Retrieved from the Internet: <URL:http://www.tr1.ibm.com/projects/xml/x ss4j/docs/xacl-ex-simple1.html> 'retrieved on 2002-04-22!	33
A	the whole document	1-34
A	HANDLY/SCHULZRINNE/SCHOOLER/ROSENBERG: "SIP: Session Initiation Protocol" IETF - RFC2543BIS-02, 'Online! 24 November 2000 (2000-11-24), XP002196891 Retrieved from the Internet: <URL:http://www.cs.columbia.edu/sip/drafts /draft-ietf-sip-rfc2543bis-02.pdf> 'retrieved on 2002-04-22!	1
A	the whole document	
A	LENNOX, J. - SCHULZRINNE, H.: "Feature Interaction in Internet Telephony" PROCEEDINGS OF THE SIXTH INTERNATIONAL WORKSHOP ON FEATURE INTERACTIONS, 'Online! 1 May 2000 (2000-05-01), XP002196892 Retrieved from the Internet: <URL:http://www.cs.columbia.edu/hgs/paper s/Lenn0005_Feature.pdf> 'retrieved on 2002-04-22!	
	the whole document	

-/--

INTERNATIONAL SEARCH REPORT

International Application No

PCT/EP 02/04867

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>BESSLER S.: "A Service Platform for Internet-Telecom Services using SIP" 6TH INTERNATIONAL CONFERENCE ON INTELLIGENCE IN NETWORKS, 'Online! 18 - 22 September 2000, XP002196893 Retrieved from the Internet: <URL:http://www.ftw.at/Dokumente/smartnet.pdf> 'retrieved on 2002-04-22! the whole document</p> <p style="text-align: center;">-----</p>	1-34



INVESTOR IN PEOPLE

Application No: GB 0224187.5
Claims searched: 1 to 14

Examiner: Andrew Courtenay
Date of search: 14 March 2003

Patents Act 1977 : Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
A, E	-	WO 2002/091756 A1 (ERICSSON)
A	-	EP 0833526 A1 (BRITISH TELECOMM)
A	-	US 6327618 B1 (CISCO)
A	-	GB 2363284 A (HEWLETT PACKARD)

Categories:

X Document indicating lack of novelty or inventive step	A Document indicating technological background and/or state of the art.
Y Document indicating lack of inventive step if combined with one or more other documents of same category.	P Document published on or after the declared priority date but before the filing date of this invention.
& Member of the same patent family	E Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^v:

H4K

Worldwide search of patent documents classified in the following areas of the IPC⁷:

H04L, H04Q

The following online and other databases have been used in the preparation of this search report:

EPODOC, WPI, JAPIO